How Can Contract Regulations Address Adverse Selection? Evidence from a Health Insurance Marketplace*

Myles Wagner†

November 1, 2024

Abstract

Rules governing contract design are a pervasive feature of regulated insurance markets. This paper asks how those rules can be used to address adverse selection. I provide evidence from health insurance, combining theory with data from Massachusetts' ACA marketplace. When cost-sharing is standardized across plans, I find severe adverse selection against broad provider network coverage, which attracts high-use patients. Alternatively, pairing broad-network coverage with *greater cost-sharing* differentially lowers demand and cost among high-use patients, weakening adverse selection. In welfare analysis, doing so yields substantial gains over standardized cost-sharing, both by preventing unraveling of broad-network coverage and improving sorting between plans. The analysis illustrates how contract regulation can complement financial transfers, like risk adjustment, in addressing adverse selection.

JEL Codes: D82, G22, I13, I18

*A previous version, which was titled "Health Insurance Menu Design with Multi-Dimensional Plans," was my job market paper in 2022-23. I am deeply indebted to my advisors David Cutler, Nathan Hendren, and Mark Shepard for their patience, guidance, and support. I am also grateful to Leila Agha, Michael Chernew, Liran Einav, Oliver Hart, Kurt Lavetti, Tim Layton, Robin Lee, Jeremy Majerovitz, Victoria Marone, Ariel Pakes, Anna Russo, Nicholas Tilipman, Benjamin Vatter, Anthony Yu, and Matt Zahn for helpful comments and discussions, as well as to many participants at the Harvard Labor/Public Finance Workshop, Harvard Industrial Organization Workshop, Harvard Department of Healthcare Policy Health Economics Seminar, NBER Aging and Health Fellows Workshop, and American Society of Health Economists meetings. This research was supported by National Institute on Aging, Grant Number T32-AG000186.

†Department of Economics, The Ohio State University. Email: wagner.1858@osu.edu

1 Introduction

One of the key promises of an insurance market is that society benefits from allowing consumers to choose among a variety of contracts provided by competing insurers. However, insurance markets may suffer from adverse selection because insurance demand typically co-varies with insurance costs (Einav, Finkelstein and Mahoney, 2021). In classic markets, adverse selection consists of plan and price distortions arising from insurers competing to attract low-cost customers (Rothschild and Stiglitz, 1976), and consumers making (socially) inefficient choices because prices don't reflect their own risk (Akerlof, 1970; Bundorf, Levin and Mahoney, 2012; Marone and Sabety, 2022).[1] These market failures, and policies for addressing them, have attracted extensive study by economists. In health insurance markets, attention has focused primarily on a standard toolbox of *financial* incentives intended to mitigate adverse selection by correcting insurer and consumer behavior. These include risk adjustment payments to insurers and subsidies and/or mandates for consumers, which are widely implemented in regulated health insurance markets (Geruso and Layton, 2017).

In this paper, I ask how *non-financial* "contract regulations" governing plan design can be used to make insurance markets robust to adverse selection. Health insurance plans are complex financial products which typically specify *cost-sharing* (how much a patient pays out-of-pocket when getting care) and *networks* (the set of hospitals and other healthcare providers that are covered), in addition to other characteristics. Health insurance regulators, in turn, often impose rules which restrict the characteristics of plans that insurers can sell in a given market. Such rules are a prominent feature of the regulatory environment in "managed competition" (Enthoven, 1988) marketplaces, such as the Affordable Care Act (ACA) exchanges and Medicare Advantage, which play a large and growing role in the provision of health insurance in the United States.[2] Despite the ubiquity of non-financial contract regulations in these settings, comparatively little work has examined their impact on – and potential utility in addressing – adverse selection.

This paper's central idea is that bundling offsetting high- and low-quality in separate plan characteristics can weaken adverse selection. An example of this in health insurance is pairing generous cost-sharing with narrow (restrictive) networks, and broad networks with restrictive cost-sharing. I call this *diagonal differentiation*. The idea follows from a multi-dimensional screening logic: if high-cost consumers have greater demand for quality of all kinds, then diagonal differentiation reduces the (adverse) correlation between overall demand and cost.

I present a simple theoretical model and empirical evidence from a health insurance marketplace to formally examine this idea. Market surplus, or welfare, depends on the efficiency of consumer sorting between plans and the price distortions resulting from insurer competition. Combining policy-driven price variation and novel administrative data, I estimate an empirical model of health insurance demand and cost in the Massachusetts ACA marketplace. My estimates imply that diagonal differenti-

[1] Pricing an individual's insurance based on their particular level of risk is not possible when that risk is private information. In practice, in health insurance markets, prices are often constrained by "uniform pricing" rules which substantially restrict the degree to which insurers can charge different prices to different consumers for the same product.

[2] In 2022, for example, more than 40 million individuals in the U.S. purchased either an ACA or Medicare Advantage health plan. An additional 30 million Americans enrolled in traditional Medicare, which is provided by the federal government, but had the option of choosing a Medicare Advantage plan.

ation substantially improves sorting and reduces competitive pricing inefficiencies, relative to a baseline of standardized cost-sharing across all plans. These findings illustrate that thoughtful contract regulations may be a powerful policy tool to address adverse selection, particularly when traditional financial incentives like risk adjustment and subsidies work imperfectly.[3] They also highlight the interaction between adverse selection and contract regulations more generally, even if they are not aimed at adverse selection. For instance, I find that standardizing cost-sharing design – the current policy in the Massachusetts ACA – may have the (presumably) unintended consequence of exacerbating adverse selection against plans with high-quality provider networks.

I begin the paper in Section 2 by presenting a stylized model of an insurance marketplace which builds on the welfare and competition framework developed by Einav, Finkelstein and Cullen (2010) and Handel, Hendel and Whinston (2015). Following that framework, I consider a two-plan market with insurers competing to provide fixed contracts which differ exogenously in the quality (breadth) of their provider networks. An important novelty in my setting is that the plans can also differ in their cost-sharing generosity, which captures the potential role of contract regulation. Patients have demand for both network breadth and cost-sharing, which are driven by their health risk, and idiosyncratic preferences for greater provider access or financial coverage. Their insurance costs under either plan also depends on their health risk, preferences, and that plan's cost-sharing. When sicker patients have greater demand for both financial coverage *and* network quality, standardizing cost-sharing makes the broad-network plan particularly attractive to high-cost sick patients.Diagonal differentiation – requiring the narrow-network plan to have more generous cost-sharing – weakens the tendency of sick patients to sort into the broad-network plan.

To build intuition for the formal results, I show how the effects of contract regulation can be illustrated graphically in the selection market framework of Einav, Finkelstein and Cullen (2010). Market surplus is determined by the equilibrium price and the area between a demand curve and a marginal cost curve. If the marginal cost curve is very steep, surplus can be small or even negative *at any price* – corresponding to the case of "backwards sorting" described by Marone and Sabety (2022). Steeper marginal cost curves also tend to generate larger pricing inefficiencies in equilibrium. Changes to contract regulation *rotate* both demand and marginal cost curves. In particular, diagonal differentiation can result in a shallow or flattened marginal cost curve by weakening the relationship between demand and cost. However, diagonal differentiation may also flatten the *demand* curve, since the plans are closer in aggregate quality. Whether overall market surplus increases from the combined effects on consumer sorting and pricing inefficiencies depends on the joint distribution of health risk, preferences, and their relationship to demand and cost – and is therefore an empirical question.

To implement this framework empirically, I turn to the Massachusetts ACA marketplace, also

[3]Financial incentives such as risk adjustment and subsidies are typically conditioned on observable individual characteristics which predict insurance costs, such as health risk (medical diagnoses), age, and geographic location. However, insurance costs can also vary with individual preferences, e.g., for particular hospitals, physicians, or prescription drugs, or one's overall taste for using healthcare. These preferences are also likely to affect plan choice, but are difficult to observe. This creates a problem of "selection on moral hazard" (Einav et al., 2013) which can lead to adverse selection even with sophisticated risk adjustment and subsidy policies. Contract regulations, such as diagonal differentiation, are complementary to these financial incentives in the sense that they operate by screening consumers on their unobserved preferences.

known as the Massachusetts Connector. In addition to being an important policy setting, the Connector is particularly well-suited to illustrate my conceptual framework because it features rich plan variation along two well-defined quality dimensions. Private insurers in the marketplace are required to offer plans at each of four "metal tiers" – Bronze, Silver, Gold, and Platinum – corresponding to increasing levels of financial generosity (i.e., lower cost-sharing). The financial parameters of these tiers (which include deductibles, maximum out-of-pocket liabilities and co-pays for different types of care) are fully standardized by Connector regulations, making all plans of the same metal tier essentially identical in their financial coverage. At the same time, Connector plans differ meaningfully in their provider networks. I focus on a particular dimension of network quality: whether the plan covers Partners Healthcare, an expensive and prestigious hospital system in Massachusetts whose importance has been shown in prior work (Ericson and Starc, 2015a; Shepard, 2022). Of the three major insurers in the exchange, only one offered a broad-network plan which covered Partners, while the others offered narrow-network plans.

I draw on linked individual demographic, plan enrollment, and medical claims data from the Massachusetts All-Payer Claims Database (APCD), and leverage two policy-driven sources of price variation to identify demand and costs for both cost-sharing and network breadth. The first occurred in 2017, when a change in the Connector's subsidy policies substantially raised the relative price of the broad-network plan. Then, in 2018, the Connector coordinated with insurers to raise the relative price of all Silver-tier plans, a practice known as "Silver-loading."[4] The price variation accompanying these policy changes, and the resulting changes in plan selection, allow me to identify demand for both financial coverage and network breadth, and its relationship to insurer costs.[5]

I use this variation to estimate a structural model of plan choice and insurer costs following closely from my conceptual framework, allowing for heterogeneity by a rich set of individual characteristics. The key dimensions of individual heterogeneity are a patient's health risk observed from medical claims, and whether the patient had a history of using Partners providers before enrolling in the Connector.[6] I find that health risk is a key driver of demand for both more generous financial coverage and broad networks. Past use of Partners is strongly associated with demand for network quality, but not financial coverage.

I then use the model to simulate adverse selection and market outcomes under different contract regulation regimes. Consistent with the predictions of the conceptual model, adverse selection is severe when the regulator standardizes cost-sharing, while diagonal differentiation substantially weakens adverse selection. In my baseline analysis, diagonal differentiation between Silver-tier narrow-network and Bronze-tier broad-network plans yields about $360 per patient-year (roughly 7.7% of average

[4]Silver-loading was a response to a surprise 2018 cut in federal *cost-sharing* subsidies for low-income enrollees. By raising the prices of Silver-tier plans in particular, the exchange allowed insurers to continue offering the same coverage to low-income patients, and recoup the lost revenue via federal *premium* subsidies, which are statutorily linked to the price of Silver plans under the ACA.

[5]To the best of my knowledge, this is one of the first papers, along with Panhans (2019), to study an ACA marketplace using individual-level enrollment and claims data, and the first to do so in Massachusetts.

[6]A unique advantage of using the Massachusetts APCD is that it enables me to observe detailed medical utilization histories at the patient level, including before the patient ever enrolled in a Connector plan. I construct a measure of a patient's loyalty to Partners based on their history of using Partners providers before enrolling in the Connector, allowing me to avoid common challenges of distinguishing preference heterogeneity from state dependence.

annual costs) greater surplus than a standardized menu requiring both plans to be offered at the Silver-tier. About 44% of the relative gains ($160 per patient-year) come from more efficient consumer sorting; even when prices are efficient, diagonal differentiation does better at *screening* patients into (socially efficient) coverage. The remainder of the gains come from more efficient (less distorted) prices in competitive equilibrium under diagonal differentiation. With realistic risk adjustment, equilibrium prices in the diagonal menu are close to the socially efficient level. However, standardizing cost-sharing leads the broad-network plan to substantially unravel due to selection on heterogeneous incremental cost or "moral hazard" (Einav et al., 2013). I show that these qualitative patterns hold across a variety of alternative robustness specifications.

This paper's conceptual and empirical results are important for several reasons. First, they highlight that adverse selection may be influenced by regulations which govern how plans are allowed to differentiate, rather than being an inevitable feature of insurance markets. Adverse selection has typically been viewed as an externality which can be addressed using financial incentives like risk adjustment or subsidies. Contract regulations, while widespread in health insurance markets, have not typically been viewed as a complementary policy margin which can also address adverse selection.[7] This paper describes how those rules can interact with key selection challenges, and may provide guidance to beneficial market design policy.

Second, the ideas in this paper may help explain observed patterns in real-world insurance settings. For example, previous work has noted the prevalence of narrow-network plans in the ACA exchanges, for which adverse selection is a possible explanation (Shepard, 2022). Severe adverse selection against broad-network plans in the ACA may be at least partially the result of standardization policies that explicitly prevent those broad-network plans from differentiating on financial coverage. This is analogous to the insight of Handel (2013) that choice architecture can exacerbate adverse selection. In the case of plan standardization, I highlight an additional mechanism, which is that consumer sorting may be inefficient even when prices are efficient. Anecdotally, diagonal differentiation, e.g., choice between a low cost-sharing HMO (narrow-network) and a high cost-sharing PPO (broad-network), are common in employer-sponsored insurance settings. My findings are consistent with prior work which has studied such menus and noted they display little adverse selection (Bundorf, Levin and Mahoney, 2012). A key contribution of this paper is to explain why this may be the case, and show formally when such menus improve welfare.

Broadly, this paper relates to several strands of literature in the economics of selection markets. It complements a large literature on policy trade-offs in regulated insurance markets under "managed competition" (Enthoven, 1988), including papers studying demand for health insurance in individual exchanges (e.g., Dafny, Ho and Varela, 2013; Dafny, Hendel and Wilson, 2015; Ericson and Starc, 2015a, 2016), and policies designed to promote efficiency in health plan provision and pricing under

[7] Menu design may be well-suited to address certain selection problems which pose challenges for financial incentive policies. For example, risk adjustment policies may have difficulty addressing adverse selection in cases where selection and incremental costs ("moral hazard") are driven by difficult-to-observe characteristics like demand for particular provider services or prescription drugs (Carey, 2017; Lavetti and Simon, 2018; Geruso et al., 2019; Shepard, 2022). Designing a menu where a plan which covers more expensive providers or drugs also charges higher cost sharing (diagonal differentiation) automatically compensates it for enrolling patients particularly likely to use those providers or drugs, in a way that is difficult for traditional risk adjustment to achieve.

competition (e.g., Einav et al., 2010; Handel et al., 2015; Azevedo and Gottlieb, 2017; Geruso et al., 2021). It also relates to the literature on the effects of financial (Einav et al., 2013) and provider network (Shepard, 2022) health plan characteristics on selection and costs, and follows Finkelstein and Poterba (2004) and Finkelstein and McGarry (2006) in emphasizing the importance of multi-dimensional preferences and costs in selection markets.

Most directly, this paper contributes to a growing literature in the design and pricing of menus in selection markets. This includes Bundorf, Levin and Mahoney (2012), Marone and Sabety (2022), Ho and Lee (2022) and Tilipman (2022) who study health insurance menu design in the employer context, and Landais et al. (2021) who study menu design in the context of unemployment insurance. This previous body of work has considered menu design primarily as a *demand-side* intervention to improve consumer screening in settings where plans are supplied and prices set by a central regulator. This paper additionally highlights the role that menu design regulations play as a *supply-side* policy which affects stability in competitive selection markets.

The rest of the paper proceeds as follows. Section 2 describes my conceptual framework and uses it to illustrate how menu design interacts with adverse selection. Section 3 describes the Massachusetts ACA exchange setting and the enrollment and claims data I use in estimating the empirical models of demand and cost described in Section 4. Section 5 uses these estimates to put numbers on the ideas from the conceptual framework, examining various counterfactual menus. Finally, Section 6 concludes.

2 Conceptual Framework

2.1 Setup

Following the approach of Einav, Finkelstein and Cullen (2010) and Handel, Hendel and Whinston (2015), I consider a stylized model of an insurance market with two plans, which are called $j \in \{L, H\}$. These plans differ exogenously in their non-financial quality. To make ideas concrete, and to align with my empirical setting, I refer to non-financial quality as network breadth or network quality interchangeably throughout the paper, but the framework applies equally to settings where non-financial quality involves additional dimensions such as prescription drug coverage.

Each plan's network quality is represented by the contract parameter $\lambda_j \geq 0$, which can be interpreted as an index for the per-unit cost of care covered by the plan's network (e.g., cost per office visit, procedure, etc.). It is assumed $\lambda_L < \lambda_H$ – meaning L covers a lower-quality (in the sense that it is less expensive) network than H. For instance, plan L may be a limited-network HMO plan which only covers low-cost hospitals and physicians' offices, while plan H is a broad-network fee-for-service plan which also covers high-cost hospitals and offices.

The plans may also differ in their financial coverage of medical costs, or financial generosity. Generosity is parameterized by $m_j \geq 0$, which is interpreted as an index for the out-of-pocket financial risk that an individual faces when enrolled in plan j, with greater m_j indicating a more financially generous plan.[8] For example, a Gold-tier plan under the Affordable Care Act is more financially

[8] While financial coverage may be multi-dimensional, parameterizing financial generosity along a single dimension m can capture any setting in which financial features are restricted in such a way that, between any two possible plans, one

generous than a Silver-tier plan. Unlike networks, which are taken to be exogenous, a market regulator can exercise menu design through the respective financial generosity levels of L and H. Individuals participating in the market purchase exactly one plan, according to preferences described below.

Consumer Preferences Individuals, indexed by i, have multiple dimensions of heterogeneity. The first dimension is an individual's health risk, $r_i \geq 0$, which represents a *baseline quantity* of healthcare individual i expects to consume while enrolled; patients with greater r_i are said to be sicker (or riskier) than those with lower health risk. Individuals may also vary in their taste for network quality or financial coverage conditional on their health risk. These preferences are represented by the parameters θ_i^λ (taste for network quality) and θ_i^m (taste for financial coverage). Individuals have quasi-linear utility from being enrolled in plan j, which is additively separable in financial coverage and network quality, i.e., $V_{ij} = \Psi\left(m_j; r_i, \theta_i^m\right) + \Phi\left(\lambda_j; r_i, \theta_i^\lambda\right)$.

I assume risk and preference parameters enter plan demand according to the following linear functional form,

$$V_{ij} = \underbrace{\left(\theta_i^m + \alpha r_i\right) \cdot m_j}_{\text{Demand for financial coverage}} + \underbrace{\left(\theta_i^\lambda + \beta r_i\right) \cdot \lambda_j}_{\text{Demand for network quality}} \tag{1}$$

Note that $\Psi_{m,i} \equiv \frac{\partial \Psi_i}{\partial m}$ and $\Phi_{\lambda,i} \equiv \frac{\partial \Phi_i}{\partial \lambda}$ are assumed constant, an assumption which holds for a local approximation to any differentiable utility specification. Demand for financial coverage and network quality can each be further decomposed, as shown in Equation (1), into *risk-driven* and *preference-driven* parts. To see this, consider the linear projection $\Phi_{\lambda,i} = b + \beta r_i + \varepsilon_i$, where $\beta = \frac{\text{Cov}(\Phi_\lambda, r)}{\text{Var}(r)}$ is the regression coefficient, and ε is mean-zero and uncorrelated with r. By construction, $\theta_i^\lambda = b + \varepsilon_i$ captures heterogeneity in demand for network quality, conditional on health risk – that is, $\text{Cov}\left(r, \theta^\lambda\right) = 0$. A similar decomposition can be performed for demand for financial coverage Ψ_m.

In order to simplify exposition throughout the discussion which follows, I assume that α and β are strictly positive, that $\theta^m = 0$ (i.e., health risk fully explains heterogeneity in demand for financial coverage), and finally that θ^λ is independently distributed from r. This set of assumptions captures the motivating intuition that sicker patients, because they make greater use of any benefits, value generous financial and network coverage more highly *on average* than healthier patients.[9]

As in Handel, Hendel and Whinston (2015), I abstract away from market participation (the extensive margin), e.g., by assuming the market regulator uses menu-invariant or budget-neutral subsidies and/or mandates to ensure individuals always participate. Choice is therefore only between L and H, and is determined by the relative price, $\Delta p = p_H - p_L$. To capture adverse selection, all individuals pay the same prices. Because demand is quasi-linear, individual i chooses H if and only if $\Delta V_i \equiv V_{iH} - V_{iL} \geq \Delta p$. Plan selection is thus determined by *incremental* demand,

$$\Delta V_i = \underbrace{r_i \cdot \left[\alpha\left(m_H - m_L\right) + \beta\left(\lambda_H - \lambda_L\right)\right]}_{\text{Risk-driven selection}} + \underbrace{\theta_i^\lambda\left(\lambda_H - \lambda_L\right)}_{\text{Preference-driven selection}} \tag{2}$$

plan first-order stochastically dominates the other in terms of a patient's distribution of out-of-pocket expenses.

[9] These assumptions are for expositional convenience, as they illustrate the key intuitions and streamline interpretation of the conceptual results. I do not impose them in the empirical analysis.

Equation (2) decomposes plan choice into a component associated with health risk alone, and another associated solely with preference heterogeneity. Selection on health risk and selection on preferences may respectively lead to different patterns of selection on *costs* (adverse selection), which has implications both for the efficiency of consumer choice and plan provision. I discuss insurer costs below.

Insurer Costs An insurer's cost of providing coverage to a given individual depends on plan characteristics (m_j, λ_j), and varies across individuals. Consumer type $\left(r_i, \theta_i^\lambda\right)$ is sufficient to describe choice behavior, so it is necessary only to specify the expected cost of covering individuals conditional on their type, e.g., $C_{ij}\left(m_j, \lambda_j; r_i, \theta_i^\lambda\right)$. For a given individual, higher-quality networks and greater financial coverage each cost more to provide. Mechanically, a more financially generous insurer pays a greater share of any medical costs (actuarial value) on average. Similarly, a broader-network plan which covers more expensive providers will pay a higher (average) price for any care delivered. I require that the cost of higher-quality coverage is increasing in a patient's health risk, r_i, to reflect that these costs accrue on a per-unit-of-healthcare basis. The following simple specification captures this basic intuition,

$$C_{ij} = \underbrace{r_i \cdot m_j \cdot \lambda_j}_{\text{Average cost proportional to } r_i} + \underbrace{k\theta_i^\lambda \cdot (\lambda_j - \lambda_L)}_{\text{Preference-driven heterogeneity}} \tag{3}$$

where r_i is defined as a patient's expected quantity of healthcare consumption *using plan L as a baseline*, that is, individual i costs $r_i \cdot m_L \cdot \lambda_L$ in expectation to cover on plan L. Importantly, Equation (3) also allows for *preference-driven heterogeneity* in the cost effect of changing plan features relative to the baseline, through the second term, $k\theta_i^\lambda \cdot (\lambda_j - \lambda_L)$. Conditional on an individual's health risk in the baseline plan, r_i, cost under a higher-quality network may vary by network preference type θ_i due to "selection on moral hazard" (Einav et al., 2013). Patients who are especially likely to use expensive providers may be expected to have greater demand for a high network-quality plan which covers those providers (Shepard, 2022), which would correspond to $k > 0$. For any plans L and H, the *incremental* cost, $\Delta C_i \equiv C_{iH} - C_{iL}$ is given by,

$$\Delta C_i = \underbrace{r_i \cdot (m_H \lambda_H - m_L \lambda_L)}_{\text{Risk-predicted incremental cost}} + \underbrace{k\theta_i^\lambda \cdot (\lambda_H - \lambda_L)}_{\text{Preference-driven heterogeneity}} \tag{4}$$

Menu Design and Diagonal Differentiation To illustrate the impact of menu design, I analyze two menu configurations the market regulator could employ. Both menus include the same "baseline" plan L with features m_L and λ_L. The regulator can offer a second plan H, where λ_H is normalized such that $\lambda_H - \lambda_L = 1$. The menus differ solely in the financial coverage level, m_H, of the second plan.

1. ***Vertical differentiation***: $m_H^{Vert} = m_L$, i.e., plans differ only in their network quality. This is analogous to prior work analyzing settings with a single dimension of plan quality (Einav et al., 2013; Azevedo and Gottlieb, 2017; Marone and Sabety, 2022). Under vertical differentiation, plan choice is determined by a combination of health risk and plan preferences,

$$\Delta V_i^{Vert} = r_i \beta + \theta_i^\lambda \tag{5}$$

2. ***Diagonal differentiation***: $m_H^{Diag} = m_L - \frac{\beta}{\alpha}(\lambda_H - \lambda_L)$. Since $\lambda_H > \lambda_L$ and $\alpha, \beta > 0$ by assumption, $m_H^{Diag} < m_L$ – i.e., plan H has a broader network but *less generous* financial coverage. That is, the diagonal menu offers choice between a plan L which has a narrow network and more generous financial coverage, and another plan H with a broad network but less generous financial coverage. The diagonal menu is defined such that the risk selection term in Equation (2) equals zero – plan choice driven solely by preferences θ_i^λ. In particular, the normalization that $\lambda_H - \lambda_L = 1$ yields that incremental demand between diagonally differentiated plans is,

$$\Delta V_i^{Diag} = \theta_i^\lambda \tag{6}$$

Next, I use this framework to illustrate the implications of menu design for the efficiency of consumer choice between plans (welfare) and plan pricing under perfect competition.

2.2 Menu Design, Choice Efficiency, and Welfare

Adverse selection distorts consumer choices between plans, since individuals have different marginal costs, ΔC_i, of switching plans, but pay the same price, Δp, to do so. As a result, consumer choice may result in small (or even negative) social gains at any price. In this subsection, I describe how menu design affects this facet of adverse selection. In order to do so formally, I introduce a statistic called *choice efficiency*, which captures how strongly an individual's *demand* for plan H is associated with society's *marginal benefit* from that individual buying plan H. I show how choice efficiency is related to the social benefit of choice in the market, and use it to examine the impact of different menu designs.

To begin, I define social (market) welfare as the total surplus, $SW \equiv E_i\left[V_{ij} - C_{ij} \,|\, i \text{ chooses } j\right]$, of individual plan choices.[10] Let $s(\Delta p)$ be the share of individuals who purchase H at price Δp. I denote an individual's marginal contribution to social welfare by choosing H relative to L, as

$$\gamma_i \equiv \Delta V_i - \Delta C_i \tag{7}$$

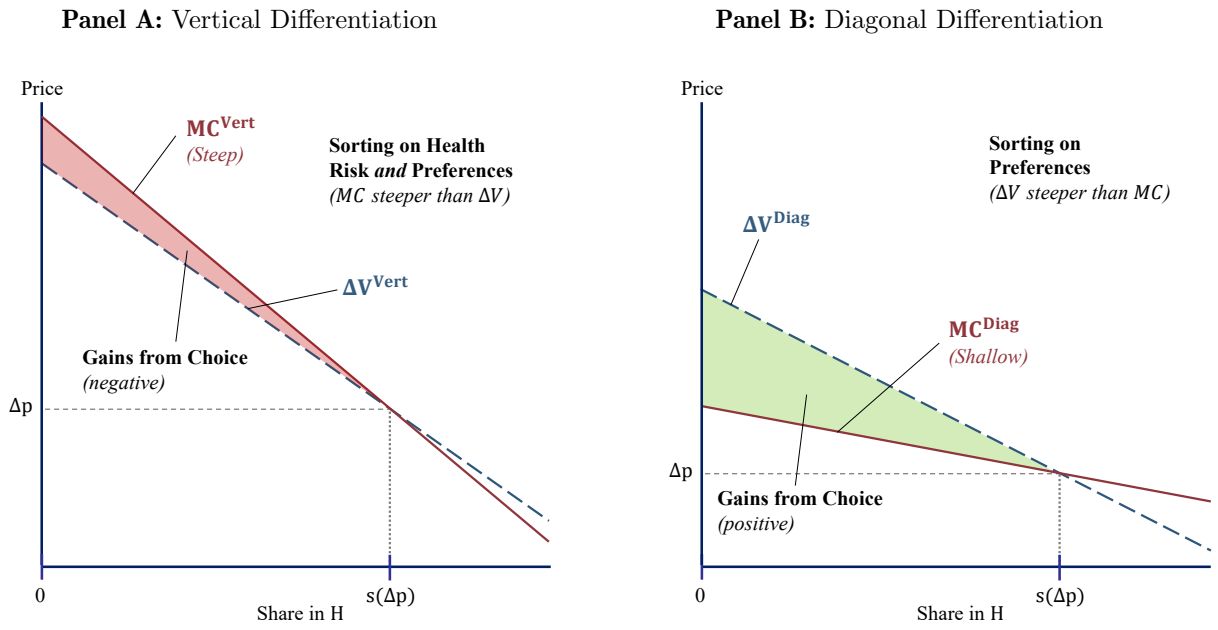
Following Landais, Hendren and Spinnewijn (2022), I refer to γ_i as the *willingness-to-pay (WTP) markup* (or simply, *markup*) individual i is willing to pay for H over L, over their own incremental cost. Market welfare would be maximized if individual i were to buy H if and only if $\gamma_i \geq 0$. However, adverse selection (uniform pricing) means that individuals instead choose H whenever $\Delta V_i \geq \Delta p$.

Let SW_L refer to the social value of enrolling all individuals in plan L, and let SW_H be the value of enrolling everyone in H. Let $\overline{\gamma} = E[\gamma] = SW_H - SW_L$ be the average markup in the population, and $\overline{\gamma}_1(\Delta p) = E\left[\gamma_i \,|\, \Delta V_i \geq \Delta p\right]$, and $\overline{\gamma}_0(\Delta p) = E\left[\gamma_i \,|\, \Delta V_i < \Delta p\right]$ denote the average markups of those who buy H and L, respectively, at price Δp. Market surplus can thus be expressed as

$$SW(\Delta p) = SW_L + \underbrace{s(\Delta p) \cdot \overline{\gamma}_1(\Delta p)}_{\text{Gains from Choice}} \tag{8}$$

[10] Note that this interprets insurer's cost C_{ij} as the social-welfare-relevant marginal cost of providing individual i with plan j. It is straightforward to amend the framework to allow for a wedge between insurer and social costs.

Figure 1: Menu Design and Choice Efficiency



Notes: The figure illustrates examples of sorting efficiency and market surplus under alternative menu designs. The dashed-blue curve corresponds to incremental demand, ΔV, and the solid-red curve plots average incremental cost $MC = E[\Delta C|\Delta V]$. The gains from choice defined in Equation (8) are given by the shaded area between ΔV and MC at relative price Δp. Panel A shows a hypothetical case for a vertically differentiated menu, where plans L and H differ only in their network quality. In the example, health risk is strongly associated with incremental costs ΔC. Consequently, the MC curve is *steeper* than the ΔV curve, resulting in negative gains from choice. Panel B shows sorting efficiency in a diagonally differentiated menu, where patients sort only on preferences by not on health risk. In the example, preferences are weakly associated with incremental cost, so ΔV is steeper than MC and choice yields positive gains.

The gains from consumer choice can be shown visually using the graphical welfare framework developed by Einav, Finkelstein and Cullen (2010). The market surplus from choice at price Δp is given by the area between two curves: the incremental *demand curve* defined by tracing out the share of individuals with $\Delta V_i \geq \Delta p$ at different prices, and the incremental marginal cost curve $MC(\Delta V) \equiv E[\Delta C \,|\, \Delta V]$ defined as the mean of ΔC_i for individuals at a given level of incremental demand. Figure 1 illustrates a pair of examples, corresponding to alternate menu designs.

The Choice Efficiency Heuristic Graphically, the sign and magnitude of the gains from choice are related to the *relative steepness* of the demand (ΔV) and marginal cost (MC) curves. Panel A of Figure 1 shows a case where the MC^{Vert} curve is steeper than the demand (ΔV^{Vert}) curve. As a result, the gains from choice are negative (shaded in red): those with the highest willingness-to-pay for H have negative average markups γ. This is an example of "backward sorting" (Marone and Sabety, 2022), in which the market regulator is better off eliminating either L or H from the choice set (effectively shutting down the market) than offering choice *at any price*. On the other hand, Panel B of Figure 1 shows the case for diagonal differentiation. In this example, the ΔV^{Diag} curve is steeper

than MC^{Diag}, i.e., those with higher WTP for H have larger and positive average values of γ, resulting in positive social gains from consumer choice.

The gains from choice are related to a heuristic I call *choice efficiency*, defined as the covariance between γ and ΔV,

$$\text{Choice Efficiency} \;=\; \text{Cov}\left(\gamma\,,\,\Delta V\right) \tag{9}$$

Choice efficiency captures how strongly incremental demand, ΔV, screens on socially-efficient markups, γ. By applying the definition of γ in Equation (7), choice efficiency can be equivalently expressed as $\text{Var}\left(\Delta V\right) - \text{Cov}\left(\Delta V, \Delta C\right)$. Since greater $\text{Var}\left(\Delta V\right)$ implies a steeper ΔV curve (stronger selection on ΔV), and a larger $\text{Cov}\left(\Delta V, \Delta C\right)$ implies a steeper MC curve (stronger selection on ΔC), choice efficiency reflects the relative steepness of the ΔV and MC curves.

It is possible to precisely express the relationship between social welfare and choice efficiency. Consider the linear projection of γ onto ΔV given by $\gamma_i = \tilde{\gamma}_i + \frac{\text{Cov}(\gamma,\Delta V)}{\text{Var}(\Delta V)} \cdot \Delta V_i$. Equation (8) can be rearranged to show

$$SW\left(\Delta p\right) = \underbrace{SW_L + s\left(\Delta p\right)\cdot\overline{\gamma}}_{\text{Value of } \textit{random} \text{ sorting}} \;+\; \delta V\left(\Delta p\right)\cdot\underbrace{\text{Cov}\left(\gamma\,,\,\Delta V\right)}_{\text{Choice Efficiency}} \;+\; \delta\tilde{\gamma}\left(\Delta p\right) \tag{10}$$

where $0 \leq \delta V\left(\Delta p\right) = \frac{s(1-s)}{\text{Var}(\Delta V)} \cdot \left(E\left[\Delta V|\Delta V \geq \Delta p\right] - E\left[\Delta V|\Delta V < \Delta p\right]\right)$.[11] The final term, $\delta\tilde{\gamma}\left(\Delta p\right) = s\left(1-s\right)\cdot\left(E\left[\tilde{\gamma}|\Delta V \geq \Delta p\right] - E\left(\tilde{\gamma}|\Delta V < \Delta p\right)\right)$, allows for non-linearity in the statistical relationship between γ and ΔV; if $E\left[\gamma|\Delta V\right]$ is linear, then $\delta\tilde{\gamma} = 0$. Appendix A.1 gives a detailed derivation.

The first term in Equation (10) is also equal to $s \cdot SW_H + (1-s)\cdot SW_L$, which is the social value of randomly enrolling share s of individuals in H and share $1-s$ in L. Note that random sorting is (weakly) worse than enrolling everyone in either H or L, since one of the two plans will be (on average weakly) more socially efficient than the other. Choice can still improve welfare, but only if sorting is *sufficiently better than random*, which may be true if choice efficiency (in the second term) is sufficiently great. Choice efficiency therefore summarizes the idea illustrated in Figure (1), and is a particularly useful tool to analyze menu design because it does not depend on plan pricing. I now use it to compare the vertically and diagonally differentiated menus introduced above.

Menu Design and Choice Efficiency Recall that the diagonal and vertical menus result in different values of ΔV, given in Equations (6) and (5), respectively. Similarly, the menu affects the incremental cost, ΔC, directly through the difference in plan features. Plugging the menu definitions into Equation (4) and computing choice efficiency gives that, for a vertically differentiated menu, choice efficiency is

[11] The term $\delta V\left(\Delta p\right)$ is a scaling factor which depends on the *shape* of the distribution of ΔV and the share of individuals, $s\left(\Delta p\right)$, who choose H at price Δp. Intuitively, the social benefits from choice are small if $s\left(\Delta p\right)$ is close to zero or one. It may be possible to express $\delta V\left(\Delta p\right)$ more precisely by making a distributional assumption on ΔV; in general it can be bounded. Note that if $\sigma^2_{\Delta V}$ denotes the variance of ΔV, Chebyshev's inequality implies,

$$E\left[\Delta V|\Delta V \geq \Delta p\right] - E\left[\Delta V|\Delta V < \Delta p\right] \leq \frac{\sigma_{\Delta V}}{\sqrt{s(1-s)}},$$

which then gives $0 \leq \delta V(\Delta p) \leq \frac{\sqrt{s(1-s)}}{\sigma_{\Delta V}}$.

given by,

$$\text{Cov}\left(\gamma^{Vert}, \Delta V^{Vert}\right) = \underbrace{\beta\left(\beta - 1\right)\sigma_r^2}_{\text{Efficiency of sorting on health risk}} + \underbrace{\left(1-k\right)\sigma_{\theta^\lambda}^2}_{\text{Efficiency of sorting on preferences}} \tag{11}$$

where σ_r^2 and $\sigma_{\theta^\lambda}^2$ denote the variance of r and θ^λ, respectively. On the other hand, choice efficiency in the diagonally differentiated menu is given by,

$$\text{Cov}\left(\gamma^{Diag}, \Delta V^{Diag}\right) = \underbrace{\left(1-k\right)\sigma_{\theta^\lambda}^2}_{\text{Efficiency of sorting on preferences}} \tag{12}$$

Comparing Equations (11) and (12) yields the key conceptual insight of the impact of menu design on choice efficiency. When $k < 1$, preferences θ^λ drive willingness-to-pay more strongly than they drive cost heterogeneity, and *sorting on preferences* is efficient. However, if $\beta < 1$ and σ_r^2 is sufficiently large, a vertically differentiated menu may lead to small overall gains or even backward sorting because it also leads to *sorting on health risk* which is inefficient. This concern is likely to be empirically relevant, since heterogeneity in health risk is substantial in many health insurance settings (σ_r^2 is large), and may outweigh heterogeneity in willingness-to-pay ($\beta = \frac{\text{Cov}(\Phi_\lambda, r)}{\sigma_r^2}$ may be small).

When riskier (sicker) patients have greater demand for both network quality and financial coverage, a diagonally-differentiated menu can reduce (or, as in the presented case, completely eliminate) sorting on health risk, leaving the socially efficient sorting on preferences. By requiring the high-quality network plan (H) to have less financial generosity, the diagonal menu screens out patients who would otherwise choose H solely because they had high health risk, which is efficient if those patients have negative markups on average. Using the graphical framework of Einav et al. (2010), menu design *rotates* (and possibly shifts) both the demand (ΔV) and marginal cost (MC) curves. In the case where sorting on preferences is efficient but sorting on health risk is not, going from the vertically differentiated menu (Panel A of Figure 1) to a diagonally differentiated menu (Panel B of Figure 1), flattens both the MC curve and the ΔV curve, since variation in health risk contributes to both ΔV and ΔC. However, because health risk drives cost heterogeneity more strongly than demand heterogeneity, the MC curve is flattened further, resulting in positive gains from choice in the diagonal menu.

2.3 Menu Design, Insurer Competition, and Pricing

Section 2.2 above shows that menu design can affect how severely adverse selection distorts consumer choices. Adverse selection may also threaten markets by distorting the incentives for *insurers* to provide different kinds of coverage. Because consumers pay the same price regardless of their cost, costlier patients are less profitable to enroll, creating a disincentive for insurers to offer plans which are more attractive to costly patients. Adverse selection can therefore lead those plans to be priced higher than society would prefer, or in the extreme prevent such plans from being offered at all. In this section, I discuss how menu design affects this aspect of adverse selection.

I consider a setting in which single-plan insurers offer one of either L or H, according to the menu designed by the market regulator. The price Δp is set competitively, resulting in an equilibrium where

insurers make zero profits. Any equilibrium price must satisfy the break-even condition,

$$\Delta AC\left(\Delta p\right) = \Delta p, \tag{13}$$

where $\Delta AC\left(\Delta p\right) = E\left[C_{iH} \mid \Delta V_i \geq \Delta p\right] - E\left[C_{iL} \mid \Delta V_i < \Delta p\right]$ is the difference in average insurer costs between plans H and L. Handel, Hendel and Whinston (2015) show conditions under which a unique Riley equilibrium always exists which satisfies Equation (13), which can be verified for the setup above (assuming a continuous distribution of consumer types).[12] Since $C_{iH} = C_{iL} + \Delta C_i$, and $C_{iL} = r_i \cdot m_L \cdot \lambda_L$,the difference in average insurer costs can be expressed as

$$\Delta AC\left(\Delta p\right) = \; m_L\lambda_L \cdot \underbrace{\left(E\left[r_i \mid \Delta V_i \geq \Delta p\right] - E\left[r_i \mid \Delta V_i < \Delta p\right]\right)}_{\text{Sorting by health risk}} \; + \; \underbrace{E\left[\Delta C_i \mid \Delta V_i \geq \Delta p\right]}_{\text{Selection on incremental cost } (\Delta C_i)} \tag{14}$$

which is the sum of two selection forces: one is how strongly consumers sort between L and H on their baseline cost (related to their health risk), and the other is how strongly consumers select on their incremental costs (which depends both on their health risk and their plan preferences θ_i^{λ}). The impact of adverse selection is more severe when $\Delta AC\left(\Delta p\right)$ is large and increasing, because the equilibrium price may be very high. Next, I show how menu design affects $\Delta AC\left(\Delta p\right)$, and how it interacts with risk adjustment, which is a standard incentive-based policy designed to address adverse selection.

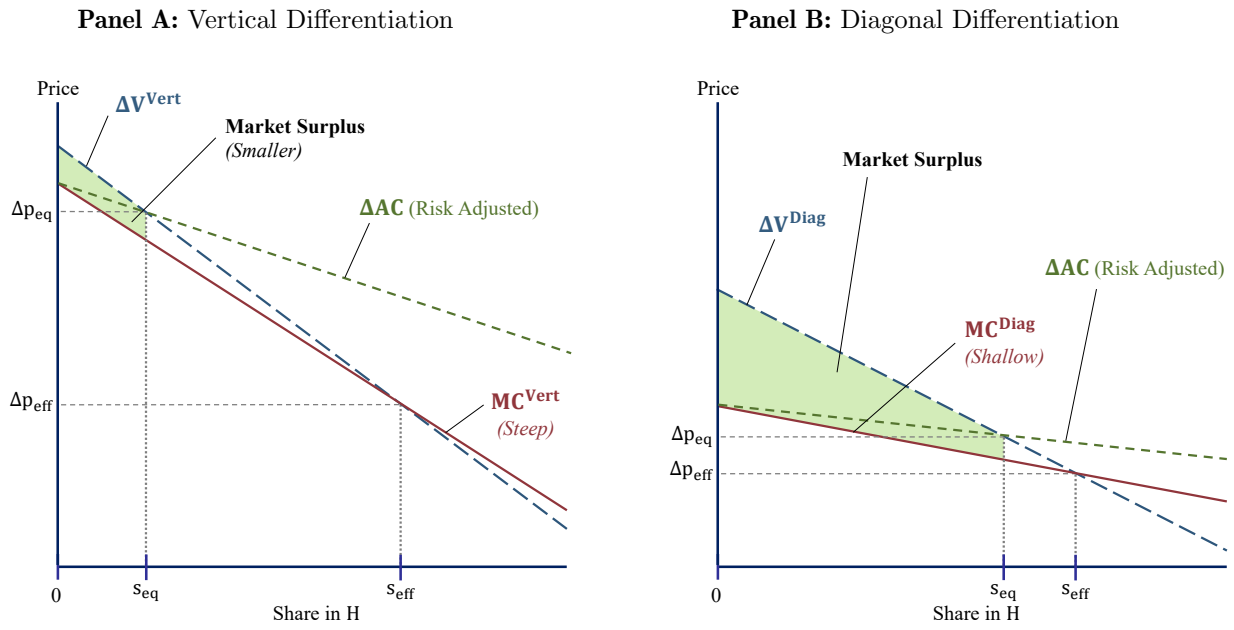
Risk Adjustment When sicker patients have greater demand for H, the first term in Equation (14) is positive, and may be very large at any price Δp, because patients sort on their health risk. Risk adjustment is a policy tool designed to address this problem, by directly compenstating insurers who provide H according to the health risk of the patients they enroll. Mechanically, a market regulator implements risk adjustment by paying each insurer a transfer τ_i for each patient enrolled. This generates a risk-adjusted difference in average costs experienced by the insurers,

$$\Delta AC^{\text{Risk Adj}}\left(\Delta p\right) = \Delta AC\left(\Delta p\right) - \underbrace{\left(E\left[\tau_i \mid \Delta V_i \geq \Delta p\right] - E\left[\tau_i \mid \Delta V_i < \Delta p\right]\right)}_{\text{Difference in risk adjustment transfer}} \tag{15}$$

which determines equilibrium pricing. The transfer τ_i is typically based on a prediction, $\hat{r}_i$, of each patient's health risk from (either ex-ante or ex-post) observable diagnoses and procedures, scaled or calibrated to cost in a baseline plan. I consider the case in which the regulator uses plan L as a baseline, that is the risk adjustment transfer is, $\tau_i = m_L\lambda_L \cdot \hat{r}_i$.

[12] In particular, their proof requires that $E\left[C_{iH} \mid \Delta V_i \geq \Delta p\right]$ and $E\left[C_{iL} \mid \Delta V_i < \Delta p\right]$ are strictly increasing in Δp.

Figure 2: Menu Design and Market Equilibrium



Notes: The figures illustrate market equilibrium with risk adjustment under different menu designs. Equilibrium is given by the break-even point, where the risk-adjusted ΔAC curve (short-dashed green curve) intersects with the demand ΔV curve (long-dashed blue curve). Both panels illustrate cases where the risk adjustment payment for individual i is given by $\tau_i = m_L \lambda_L \cdot \hat{r}_i$, so that the ΔAC curve is equal to the integral of the marginal cost curve ΔC (solid red curve). Panel A illustrates an example for a vertically differentiated menu, where adverse selection is severe because the marginal cost curve is steep. In Panel B, the marginal cost curve is shallow due to diagonal differentiation, so adverse selection is relatively mild.

Menu Design and Pricing with Risk Adjustment Figure 2 illustrates the case where the regulator is able to do risk adjustment based on an accurate prediction of health risk, i.e. $\hat{r}_i = r_i + \varepsilon_i$, where ε_i is classical measurement error. When this is the case, the risk-adjusted difference in average costs reduces to only the final term in Equation (14),

$$\Delta AC^{\text{Risk Adj}}(\Delta p) = \underbrace{E\left[\Delta C_i \mid \Delta V_i \geq \Delta p\right]}_{\text{Selection on incremental cost } (\Delta C_i)}$$

which is the integral of the MC curve over the portion of the distribution of ΔV_i that is greater than Δp (Einav et al. (2010) call this the AC curve). The figure illustrates the importance of the steepness of the MC curve for market surplus, even under perfect risk adjustment. Because the diagonally differentiated menu flattens the MC curve, it also flattens the risk-adjusted ΔAC curve, reducing the impact of adverse selection.

Next, consider the case of imperfect risk adjustment. Suppose health risk is measured inaccurately, e.g., $\hat{r}_i = A + Br_i + \varepsilon_i$, with $0 \leq B < 1$. Then a share $1 - B$ of variation in r_i is not compensated by risk adjustment, and there will still be some sorting on health risk in the risk-adjusted ΔAC. A diagonally differentiated menu helps here, because it reduces the amount of sorting on health risk. The

diagonal menu defined in Section 2.1 completely eliminates sorting on health risk, so it is insensitive to imperfect risk adjustment.[13] In general, reducing the importance of r_i to plan selection, reduces the *sensitivity* of ΔAC to imperfect risk adjustment.

To see this formally, note that Equation (15) becomes,

$$\Delta AC^{\text{Risk Adj}}(\Delta p) = m_L \lambda_L \cdot (1-B) \cdot \underbrace{\left(E\left[r_i \,|\, \Delta V_i \geq \Delta p\right] - E\left[r_i \,|\, \Delta V_i < \Delta p\right]\right)}_{\text{Sorting by health risk}} + \underbrace{E\left[\Delta C_i \,|\, \Delta V_i \geq \Delta p\right]}_{\text{Selection on incremental cost } (\Delta C_i)} \tag{16}$$

under imperfect risk adjustment. Relative to the vertical menu, which induces selection according to $\Delta V_i^{Vert} = \beta r_i + \theta_i^{\lambda}$, Equation (2) shows that *lowering* m_H reduces the importance of health risk to plan choice, and therefore lowers the first term of Equation (16).

3 Setting and Data

3.1 Massachusetts Connector

The Massachusetts Health Connector is the state's health insurance exchange – a regulated individual insurance market – established under the Affordable Care Act (ACA). The Connector opened in 2015, taking over for the pre-ACA Massachusetts exchange, which was established in 2006 under the state's "Romneycare" reform.[14] To the best of the author's knowledge, this paper is the first to study the Massachusetts ACA exchange using individual-level enrollment and claims data. Individuals purchasing insurance through the Connector choose from plans offered by competing private insurers. By rule under the ACA, each insurer offers a standardized menu of plans, corresponding to four "metal tiers": Bronze, Silver, Gold and Platinum. These tiers range from higher cost-sharing/less generous (Bronze), to lowest cost-sharing/most generous (Platinum).[15] The Connector strictly regulates the financial features of each of these metal tiers for insurance carriers offering plans on the exchange, but within each metal tier, provider networks differ across insurance carriers. The Connector also offers premium and cost-sharing subsidies for lower income consumers through a program called Connector-Care. Consumers falling below 300% of the federal poverty level and who are not eligible for employer insurance or other public programs (such as Medicaid) choose among generously-subsidized versions of the "baseline" Silver-tier plans, which means that, in effect, consumers below 300% of poverty choose only among insurance carriers, and not over coverage level (metal tier). Consumers above 300% of poverty do not receive state subsidies, and choose among plans that differ in coverage level (metal tier) and are offered by different insurance carriers.

The Connector setting is especially well suited to studying multi-dimensional menu design because of the rich variety of plans available. In addition to clearly defined and regulated variation in coverage level through the metal tier design, major carriers participating in the Connector differ substantially in

[13] This requires the assumption from Section (2.1) that r_i and θ_i^{λ} are independent.

[14] For research on the pre-ACA Massachusetts health insurance exchange, see e.g., Shepard (2022); Ericson and Starc (2015a,b, 2016); Hackmann et al. (2015)

[15] Insurers in the Connector also have the option of offering a least-generous "Catastrophic" tier plan; in practice, few insurers offer Catastrophic coverage and those plans that are offered see very low enrollments.

the quality (breadth) of their provider networks. I focus particularly on three major health insurance carriers in the Connector during my period of study (2015-2018): Boston Medical Center (BMC) HealthNet, Network Health (also called Tufts Health Plan Direct), and Neighborhood Health Plan (NHP). During my period of study, NHP's provider network was the only one which covered Partners Healthcare, a large system that includes Massachusetts General Hospital (MGH) and Brigham and Women's Hospital (BWH).[16] Many Partners providers, in particular MGH and BWH, are known for both their prestige and high cost of care (Ericson and Starc, 2015a; Shepard, 2022). BMC Healthnet is a vertically-integrated plan operated by Boston Medical Center, a major Boston-based public hospital which covers a narrow provider and hospital network, while Network Health is the narrow-network option offered by Tufts Health Plan, a Massachusetts-based health insurance carrier which also offers broad-network commercial networks in the employer-sponsored insurance market. The variation in network quality between BMC, Network Health, and NHP, intersected with the variation in coverage tiers mandated by the ACA, creates a discrete approximation to the two-dimensional plan space $\mathbb{P}$ described in Section 2, and provides a natural setting to study menu design with multi-dimensional plans.

3.2 Data Sources

I draw on two primary administrative data sources that, combined, contain information on plan choice sets, prices, enrollment decisions and healthcare utilization for Connector enrollees over the period 2013-2018. The first is linked enrollment and claims data from the Massachusetts All-Payer Claims Database (APCD), which I supplement with administrative enrollment and pricing data from the Connector.

Massachusetts All-Payer Claims Database (APCD) I use individual-level claims and enrollment data for the years 2013-2018 from the Massachusetts APCD, which provides linked enrollment-claims data from public and private insurance payers in the state of Massachusetts, including all carriers participating in the Connector.

The APCD enrollment data is an individual-level panel describing each record of an individual enrolled in health insurance in Massachusetts, and includes variables describing the period of enrollment, the identity of the insurance carrier, whether the plan was purchased through the individual exchange, and features of the plan such as actuarial value and ACA metal tier (Catastrophic, Bronze, Silver, Gold or Platinum). The enrollment data also describe detailed individual-level demographics including gender, age, household composition and location of residence up to a 5-digit zip code. Using individual-level identifiers, enrollment data can be linked to APCD insurance claims, which describe insurer costs, medical diagnoses and procedures, as well as the identity of hospitals and physicians providing medical care for each patient.

[16] Partners Healthcare has since re-branded and is now called Mass General Brigham, and Neighborhood Health Plan is now called Mass General Brigham Health Plan.

Massachusetts Connector Pricing Data I use administrative pricing and enrollment data for the entirety of the Massachusetts Connector over the period 2015-2018, which allow me to observe plan availability and pricing by geographic region and member age. I combine this information on plan availability and pricing to construct a dataset of individual plan enrollment decisions.

3.2.1 Plan Choice and Cost Data

I construct a comprehensive dataset at the level of each individual choice instance over plans purchased through the Connector for the period 2015-2018. Individuals in the Connector make plan enrollment decisions at two key points on the calendar. The first is when *new enrollees* initially purchase insurance through the Connector, and the second is at an annual open enrollment period during January of each calendar year, when *continuing enrollees* can switch their enrollment to a different plan. For each instance of an individual making a choice between plans, I define variables describing the set of available plans, the characteristics of these plans including the insurance carriers, metal tiers, and premiums, and the plan enrollment decisions (which plan was chosen). For each individual-choice instance, I also define patient-level variables that describe each individual's demographics, their history of healthcare use, and measures of their concurrent health (based on subsequent medical claims).

Key individual-level variables include: patient-level risk scores and total medical costs paid by the insurer (calculated over the subsequent year before the next choice instance), the average distance between the patient's zip code and Partners-affiliated hospitals, and whether the patient used Partners (defined as outpatient, non-emergency care) *before* ever enrolling in the Connector.

3.3 Price Variation in the Connector

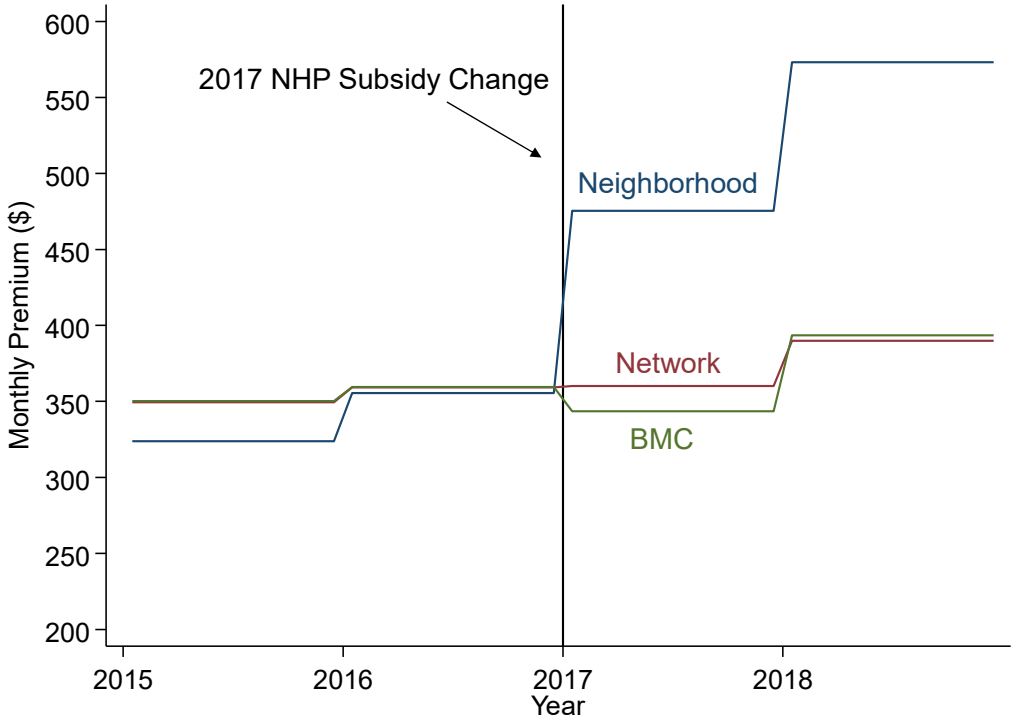
I draw on variation in plan prices over time to identify heterogeneity in plan preferences. Over the period 2015-2018, pricing by insurance carrier and by metal tier varies due to several institutional factors.

Reduced-Form Evidence: Variation in Carrier Prices

Between 2016-2017, the Connector changed the structure of its premium subsidies for plans in ConnectorCare, causing the price of NHP (the broad-network plan) to rise substantially relative to narrow-network options BMC and Network Health, and precipitating a "death spiral" (Cutler and Reber, 1998) pattern of adverse selection against NHP.[17] ConnectorCare prices are linked to each carrier's "baseline" premium in the individual market for consumers who do not receive the state (ConnectorCare) subsidies, meaning this policy change had spillover effects on the price of NHP's plans in the individual market. Figure 3 plots the average monthly premium of plans offered by each of the three major carriers over time.

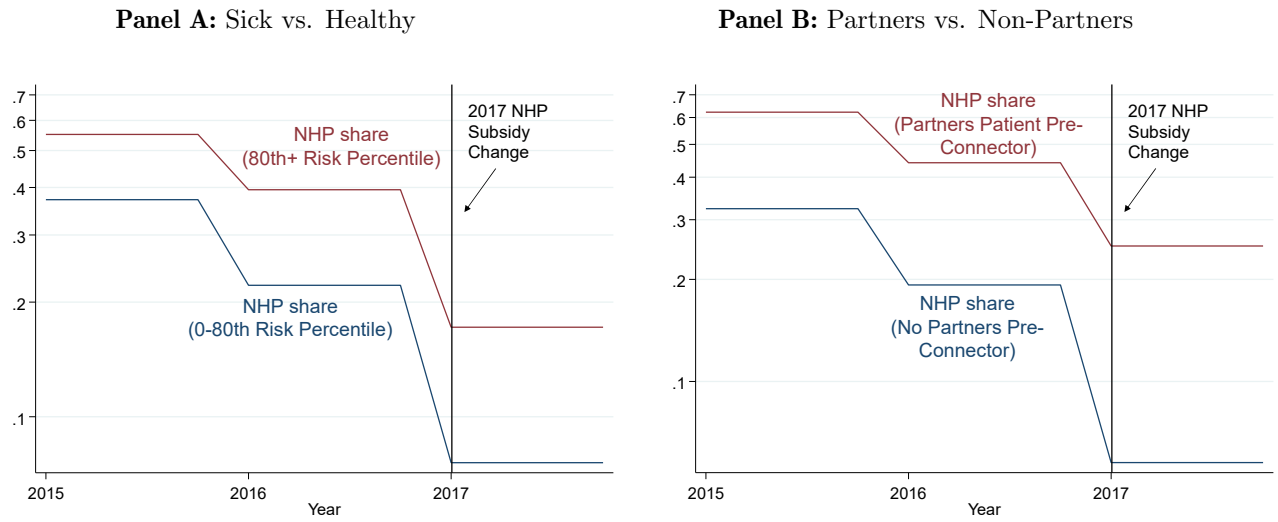
[17] Prior to 2017, the Connector granted more generous premium subsidies to NHP because of its more expensive provider network. Citing concerns that this policy created an incentive for NHP to raise its premiums, the Connector changed this policy in 2017. For research on price-linked subsidies in health insurance exchanges, see Jaffe and Shepard (2020).

Figure 3: Price Variation: Carriers



Notes: The figure shows average monthly baseline (pre-subsidy) premium by year for the plans offered in the Connector by each of the three major carriers: Neighborhood (NHP), BMC Healthnet, and Network Health, over the first four years (2015-2018) of the Connector. Averages are computed by taking the mean of the monthly premium of the lowest-price plan within each metal tier (Bronze, Silver, Gold and Platinum) offered by each of the carriers. Premiums vary by geographic region and age; averages are weighted by the distribution of individual enrollees in the Connector in 2016.

Figure 4: Heterogeneity in Demand for Neighborhood Health Plan



Notes: The figure plots market share of NHP by years 2015-2017 among new (first-time) Connector enrollees. Panel A shows heterogeneity in NHP market shares by patient risk score, and Panel B shows heterogeneity in NHP market shares by whether individual enrollees used Partners outpatient (non-emergency) care within two years before enrolling in the Connector. Shares (y-axis) are shown in log-scale. NHP shares among new enrollees in 2018 are not shown because the APCD enrollment data only describe continuing enrollees in 2018. The sample excludes all former ConnectorCare enrollees.

The change in market shares due to this price variation provides identification of the demand curve for NHP relative to BMC and Network. Figure 4 provides graphical analysis of heterogeneity in demand for NHP by sickness and history of Partners use, among first-time enrollees to the Connector. The results suggest that sickness and past use of Partners providers are both drivers of demand for the NHP plan.[18] Panel A shows that sicker patients 1) purchase NHP at greater rates, and 2) see a proportionally smaller decline in their likelihood of buying NHP plans in 2017 (when the price is higher) relative to 2016. Panel B shows a similar pattern for patients who have used Partners providers before entering the Connector, consistent with a preference for Partners coverage.

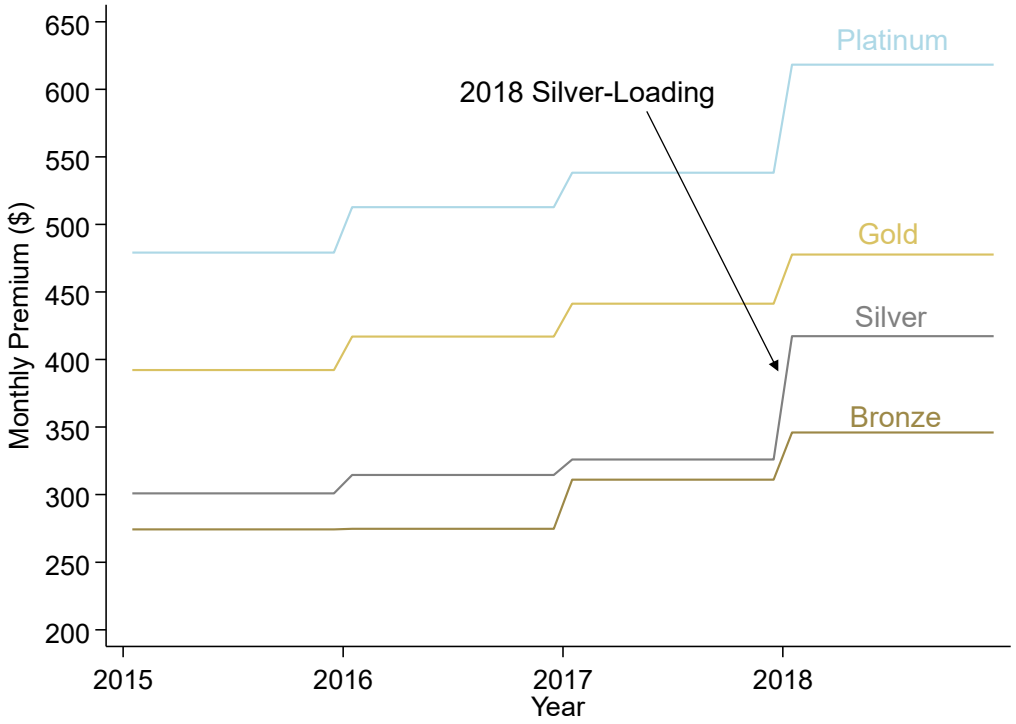
Reduced-Form Evidence: Variation in Silver-Tier Pricing

Between 2017-2018, the Connector allowed all carriers to raise the relative premiums of Silver-tier plans in order to make up for the loss of federal cost-sharing reduction subsidies, in a practice known as "Silver-loading." Figure 5 shows the variation in average monthly premiums for plans in each of the four ACA metal tiers.

Variation in the relative price of Silver versus other metal tiers identifies demand for financial coverage. Because I only observe continuing enrollees in 2018, Figure 6 presents graphical evidence of heterogeneity in demand based on the rate at which individual enrollees *switch from Silver to Bronze*

[18] The 2015-2016 price variation could in principle be used to identify another segment of the demand curve, under the assumption that unobserved carrier quality . In practice, I draw identification of demand for carrier networks from only the 2017 subsidy change, as I discuss in greater detail below in Section 4.1.

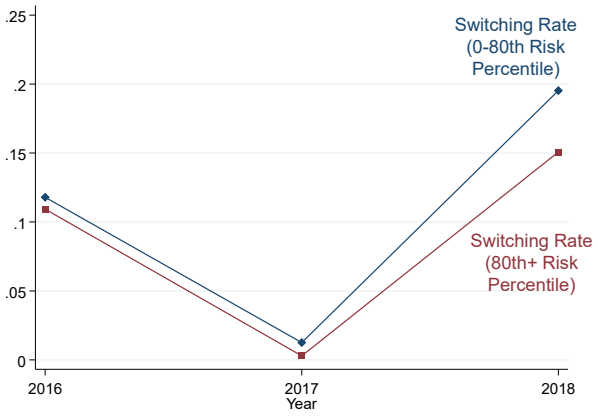
Figure 5: Silver-loading Price Variation



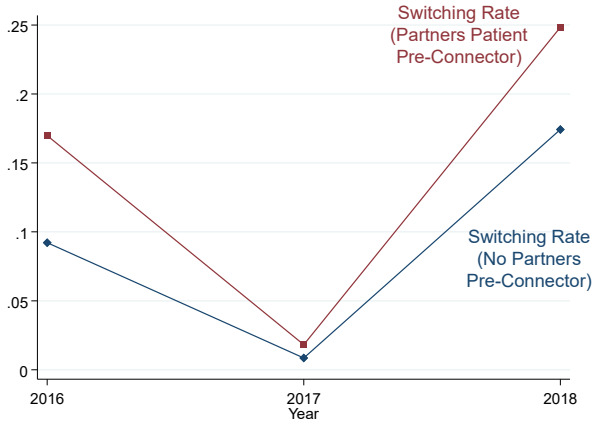
Notes: The figure shows average monthly baseline (pre-subsidy) premium by year for plans of each primary metal tier (Bronze, Silver, Gold and Platinum) offered in the Connector by each of the three major carriers: Neighborhood (NHP), BMC Healthnet, and Network Health, over the first four years (2015-2018) of the Connector. Averages are computed by taking the mean of the monthly premium of the lowest-price plan within each metal tier for each carrier, and averaging across each of the three carriers. Averages are weighted by the distribution of individual enrollees in the Connector in 2016.

Figure 6: Silver-to-Bronze Switching Rates

Panel A: Sick vs. Healthy



Panel B: Partners vs. Non-Partners



Notes: For each year y shown on the x-axis, the figure plots the rate at which enrollees in a Silver-tier plan in year $y-1$ switch to a Bronze plan during open enrollment at the start of year y. Panel A shows heterogeneity in switching rate by whether a patient is sick (top 20% of risk scores) or healthy, and Panel B shows heterogeneity in switching rate by whether a patient used Partners (outpatient, non-emergency) providers within two years before enrolling in the exchange. The sample excludes all former ConnectorCare enrollees.

coverage in response to the price change. Panel A shows that sicker patients leave the more generous Silver plan at lower rates than healthier ones when the relative price of Silver rises, consistent with sicker patients having a preference for more generous coverage.[19] Panel B shows a different pattern when switching rates are analyzed by whether the patient showed a preference for Partners providers before entering the exchange: previous Partners patients are more likely to switch from Silver to Bronze tier plans when the relative price of Silver increases.

The theory from Section 2 shows that quantifying heterogeneity in demand for plan characteristics (in the Connector setting, carrier network quality and metal tier), is a key input to the menu design problem. Given these patterns in plan choice along carrier and metal tier dimensions as a result of price variation, it is natural to ask if it is possible to quantify how much health risk, past use of Partners – and other factors such as distance to Partners hospitals, or demographics such as age – drive changes in plan shares when plans can differ in their carrier network or metal tier. Note that with sufficient price variation along each carrier-by-metal tier dimension, the demand function would be non-parametrically identified, and I could proceed in a similar fashion to Einav et al. (2010) in estimating demand curves. However, I have more limited price variation along each plan dimension *separately*, rather than independently along each carrier-by-metal tier. In order to quantify heterogeneity in demand, I therefore turn to a structural model in Section 4 and rely on parametric assumptions to bridge the gap between the price variation in my setting and the ideal price variation.

[19] An additional drawback of the APCD subset I use for this paper is that the data on medical claims for 2018 is incomplete, meaning I do not observe healthcare utilization, and cannot construct risk scores, for patients based on their 2018 claims. I therefore use the 2017 risk score for continuing patients as a proxy for 2018 health risk.

4 Empirical Model and Estimation

Section 3.3 uses price variation in the Connector to provide evidence of heterogeneity in demand for two plan features: carrier quality and metal tier, by a patient's sickness and preference for Partners healthcare providers. In Section 4.1, I present a structural model of plan demand that allows me to leverage the price variation described above into an estimate of demand for plan features that varies by a rich set of consumer observables. Section 4.2 then describes how I estimate a model of expected insurer costs, including estimating causal effects of changes to plan features on expected costs. Both the demand and cost models follow the framework developed in Section 2.

4.1 Plan Demand Model

I use the plan choice dataset to estimate a multinomial logit model of plan demand. My model takes the timing of an individual's participation in the exchange as exogenous, and models each individual's plan choice decision.[20] These decisions are made at two times: when the consumer first enters the exchange, and each year at annual open enrollment periods when enrollees are allowed to switch plans, for as long as the enrollee continues to participate in the exchange.[21] My model assumes that enrollee i choosing among plans j at time t does so in order to maximize the following utility function:

$$u_{ijt} = \underbrace{\alpha\left(z_{it}^{p}\right) p_{ijt}}_{\text{Price}} + \underbrace{\underbrace{f\left(Carrier_{jt}; z_{it}\right)}_{\text{Carrier utility (network)}} + \underbrace{g\left(MetalTier_{jt}; z_{it}\right)}_{\text{Coverage tier}}}_{\text{Empirical analog to } E[u_i(m,x)|z_{it}]} + \underbrace{\eta_{jt}\left(Age_{it}\right)}_{\text{Plan dummies}} + \underbrace{\gamma\left(z_{it}\right) 1_{ijt}^{SamePlan}}_{\text{Plan inertia}} + \epsilon_{ijt}. \tag{17}$$

In addition to an individual-plan-year-specific type-I extreme value taste shock ϵ_{ijt}, plan utility depends on prices p_{ijt}, which are observed and vary according to an individual's age and location, carrier (NHP, BMC or Network) and metal tier (Bronze, Silver, Gold or Platinum), plan inertia γ for current enrollees (those already enrolled in the Connector at the time they are making their choice), and a rich set of carrier-by-region-by-age and metal-tier-by-region-by-age fixed effects. I allow price sensitivity $\alpha\left(z_{it}\right)$ to vary by individual observables z_{it}^{p}, where z_{it}^{p} includes the enrollee's health risk and interactions of the enrollee's sex and age. I also allow carrier, metal tier and plan inertia to vary by by observables z_{it}, which includes enrollee's health risk, interactions of sex and age, interactions of distance to Partners hospitals and use of Partners providers prior to entering the exchange. While the APCD does not allow me to observe patient-level socioeconomic status, I also include measures of the patient's neighborhood (ZIP code) level SES including percent of the population that is Black, that is Hispanic, and that is below poverty in z_{it}.

[20] This follows the standard assumption in the individual exchange setting, e.g. in Shepard (2022) who studies the pre-ACA Massachusetts individual exchange. The assumption is that changes in price variation over time in the Connector affects only the decision of which plan to choose, *conditional on* participating, but does not affect the decision of whether to purchase insurance through the Connector.

[21] I exclude any enrollees who re-enter the exchange after a period of absence, due to ambiguity in the default enrollment rules for returning enrollees.

Identifying Price Variation Identification of premium coefficients requires isolated variation in prices that is orthogonal to changes in unobserved plan quality or unobserved shocks to demand. Section 3.3 describes variation in Connector prices over time, and I argue that the change in relative NHP prices from 2016-2017, and the change in relative Silver prices in 2017-2018, are each due to idiosyncratic policy shocks that are unlikely to be correlated with underlying changes in plan quality or demand. However, over the period 2015-2018, there are also other sources of price variation, which may bias demand estimates if they are correlated with unobserved plan quality. In order to ensure that my estimation uses only variation due to the aforementioned policy changes, I use a detailed set of plan dummies, $\eta_{jt}\left(Age_{it}\right)$, to absorb all variation in prices over time *except* for changes in the relative price of NHP between 2016-2017, and changes in the relative price of Silver in 2017-2018.[22]

The intuition is analogous to identification of difference-in-differences treatment effects using fixed effects. The fixed effects $\eta_{it}\left(Age_{it}\right)$ can be thought of as absorbing all variation in prices over time that may be correlated with changes in unobserved plan quality or demand (i.e., all price variation over time except that in the price of carriers between 2016-2017, and that in the price of metal tiers between 2017-2018). Because prices in the Connector vary by a consumer's age (in years) according to a fixed pricing schedule, I scale these fixed effects by this schedule as a function of age. The level of prices in the Connector also varies by geographic region, so I construct a series of region-age-scaled fixed effects. I then allow these fixed effects to vary flexibly by carrier-year and metal tier-year, with the exception of years 2016-2017 for the set of carrier fixed effects, and years 2017-2018 for the metal tier fixed effects, which are each constrained to be equal within region-age.

Another source of price variation comes from the variation in prices of carriers conditional on metal tier and variation in prices of metal tier conditional on carrier, *within* region-age-year. For example, the relative price of Silver versus Bronze may vary between NHP and BMC, or the relative price of NHP versus Network may differ between Silver and Gold within the same region-year. My identification strategy assumes this price variation is orthogonal to differences in unobserved plan quality not already captured by the fixed effects described above. This assumption seems reasonable because the Connector strictly regulates the financial features of plans in each metal tier, such that plans offered in the same tier by different carriers have the same out-of-pocket price schedules (Appendix Figures 13 and 14 give examples). Additionally, each carrier offers Bronze, Silver, Gold and Platinum plans on their standard commercial network. Instead, this variation in prices is likely due to differences in insurer *costs* due to the differences in each carrier's provider networks. In sum, the identifying assumption is that the relationship beween prices and unobserved plan quality/demand follows the structure of the fixed effects described above.[23]

4.1.1 Plan Demand Estimation and Results

I estimate this demand model via maximum likelihood on the plan choice dataset described in Section 3.2.1. Plan premiums (in dollars per month) carry a negative coefficient across all enrollee types,

[22] See e.g. Nevo (2000).

[23] Essentially, the identification assumes plan utility at each choice instance is *additively separable* in unobserved carrier and metal tier quality. Appendix C.1 discusses this assumption in more detail.

but there is substantial heterogeneity in price sensitivity: the magnitude of the premium coefficient is smaller for older and sicker individuals, which is consistent with previous estimates in a similar context (see Shepard, 2022).

The estimates imply substantial heterogeneity in demand for insurance carriers. Patients perceive BMC as the lowest quality carrier conditional on financial coverage. Although BMC and Network have similar hospital networks, Network is viewed as similar in quality to NHP by the average patient, although sicker patients and patients with a history of Partners use are willing to pay a premium for NHP coverage. Network's health plan is branded as "Tufts Direct," and owned by a large regional insurance company – Tufts Health Plan – that also offers broad-network plans (some of which cover Partners) in the employer-sponsored insurance market, which may play a role in driving Network's high valuation relative to BMC. Figure 7 plots heterogeneity in predicted willingness-to-pay (WTP) for different carriers by sickness and past Partners use, based on the distribution of enrollees in the Connector 2015-2017 (years for which I observe all new and continuing enrollees), where I define predicted WTP as

$$WTP_{ijt} = \frac{-1}{\alpha\left(z_{it}^{p}\right)} \times \left[\underbrace{f\left(Carrier_{j}; z_{it}\right) + g\left(MetalTier_{j}; z_{it}\right)}_{V_{ijt}}\right]. \tag{18}$$

Multiplying by the inverse of the coefficient on monthly premium converts the predicted utility into units of dollars per month. Note that V_{ijt} excludes plan dummies, and so holds "unobserved" plan characteristics fixed across individuals making choices at difference points in time.[24]

Figure 8 plots heterogeneity in predicted WTP for metal tiers by sickness and past Partners use. Plan valuations are increasing in coverage tier across all groups, with sickness (top 20% of risk scores) being a strong driver of demand for more generous coverage. Being a past Partners patients is also a driver of demand for more generous coverage, especially for Gold and Platinum, but a relatively weaker driver of demand for Silver relative to Bronze.[25] This pattern, with sick and past Partners patients each preferring broad-network (NHP) coverage, while sick patients have a relatively stonger preference for Silver versus Bronze, is a key aspect of demand heterogeneity that will generate heterogeneity in relative preference between plans in a diagonally differentiated menu.

4.2 Insurer Cost Model

The other input to the menu design problem is the distribution of insurer costs as a function of plan characteristics, denoted C_{ij} in Section 2. I assume the distribution of expected insurer costs follows an affine transformation of the distribution of individual-level risk scores in the population of patients in the Connector.[26]

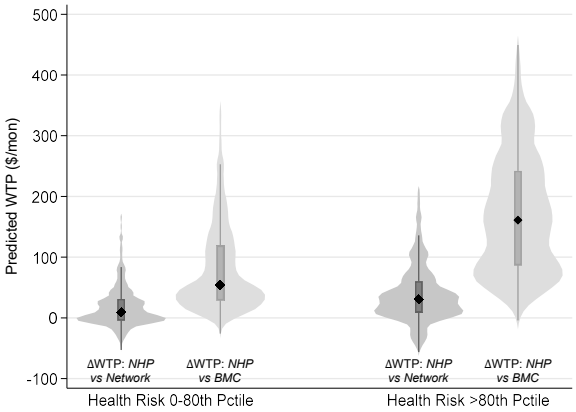
[24] I normalize plan dummies to equal zero for 2017; the predicted WTP can thus be thought of as predicted WTP for 2017 plan characteristics (carrier quality and metal tiers).

[25] Being a past Partners patient may be picking up heterogeneity in demand associated with unobserved socioeconomic factors that also drive demand for Gold and Platinum coverage, specifically.

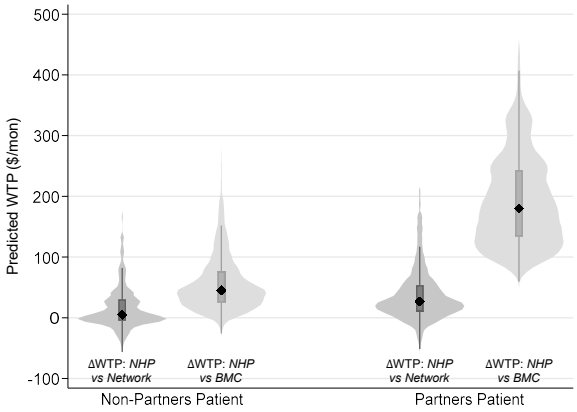
[26] These risk scores are constructed using diagnoses observed in concurrent medical claims, and represent a unitless scale measure of expected insurer costs; that is, the ratio of risk scores between any two patients approximates the ratio

Figure 7: Distribution of WTP for NHP relative to BMC and Network

Panel A: Sick vs. Healthy



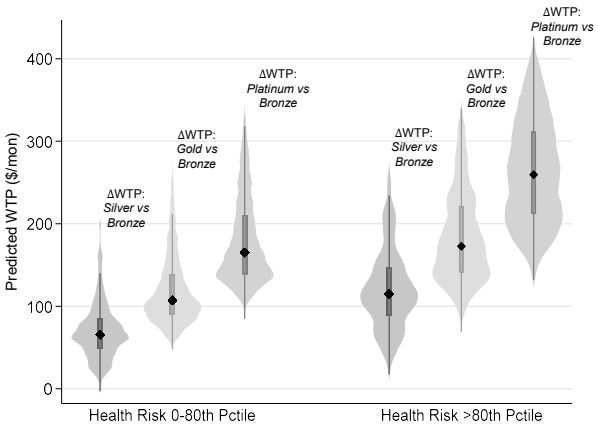
Panel B: Partners vs. Non-Partners



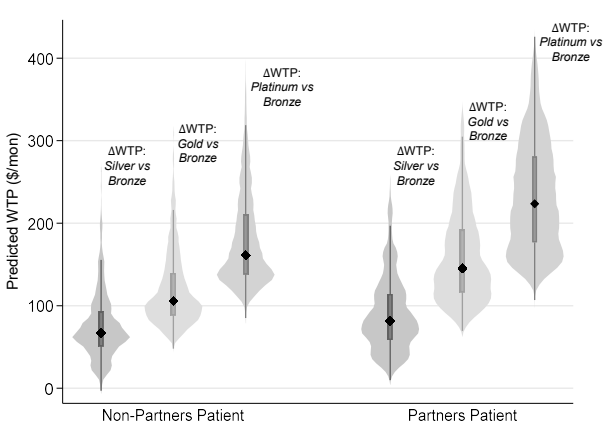
Notes: The figure plots densities of ΔWTP_{ijt} (where WTP_{ijt} is defined in Equation 18) for changes in plan carriers across various groups. Panel A compares ΔWTP for NHP versus Network and NHP versus BMC for healthy (risk score in the bottom 80%) and sick (risk score in the top 20%) patients in the Connector. Panel B compares ΔWTP for NHP versus Network and NHP versus BMC for Partners (before entering the Connector) versus non-Partners (before entering the Connector) patients. Densities are estimated on the distribution of Connector member-months 2015-2017.

Figure 8: Distribution of WTP for Metal Tiers

Panel A: Sick vs. Healthy



Panel B: Partners vs. Non-Partners



Notes: The figure plots densities of ΔWTP_{ijt} (where WTP_{ijt} is defined in Equation 18) for changes in plan metal tier across various groups. Panel A compares ΔWTP for Silver versus Bronze, Gold versus Bronze, and Platinum versus Bronze for healthy (risk score in the bottom 80%) and sick (risk score in the top 20%) patients in the Connector. Panel B compares ΔWTP for Silver versus Bronze, Gold versus Bronze, and Platinum versus Bronze for Partners (before entering the Connector) versus non-Partners (before entering the Connector) patients. Densities are estimated on the distribution of Connector member-months 2015-2017.

The other input required to define an empirical analog to C_{ij} is the causal effect of changes in plan characteristics on an insurer's expected costs of insuring a given individual. Here, I assume a proportional model of insurer costs, to reflect the idea that the incremental cost *to the insurer* of a change in coverage level or provider network is proportional to the underlying sickness of each patient. This is captured by the empirical specification

$$C_{ijt} = exp\left(\alpha_i + \beta r_{it} + \gamma_t + \delta^{covg} CS_{jt} + \lambda^{ntwk}(z_i)\, 1_{jt}^{NHP}\right), \tag{19}$$

for the insurer's expected cost of covering individual i on plan j at time t. The model includes an individual-specific (and time-invariant) individual fixed effect α_i, the individual's risk score r_{it}, and time-specific fixed effects γ_t.

The effect of an increase in coverage tier is captured by δ^{covg}; I assume that the proportional effect of an increase in coverage tier is linear in the change in cost-sharing tier CS_{jt}, that is, the proportional effect of Gold versus Silver is the same as the proportional effect of Platinum versus Gold, etc. The change in costs implied by δ^{covg} captures two separate cost effects of increased coverage generosity: the direct effect of the increase in the plan's actuarial value,[27] and "moral hazard," the increase in patient utilization due to a decrease in the out-of-pocket price of care. From the perspective of the menu design problem described in Section 2, these two channels enter the problem additively and so I treat them together.

Similarly, I assume that the broad network plan, NHP, has a proportional effect on expected insurer cost, captured by $\lambda^{ntwk}(z_i)$. Conceptually, $\lambda^{ntwk}(z_i)$ may reflect a combination of a price effect and a care intensity effect: patients on broad network plans may use a mix of providers with higher average prices, as well as providers that treat the patient *more intensively* (e.g., ordering more diagnostic tests), conditional on the patient's utilization (e.g., number of doctor's office visits). Because the cost effect of broad network coverage is likely to vary both by a patient's access to expensive/intensive providers, and the patients preference for those providers, I allow $\lambda^{ntwk}(z_i)$ to vary by observables z_i: interactions of the patient's distance from the nearest Partners hospital, and past use of outpatient Partners providers prior to enrolling in the Connector.

Estimation and Identifying Variation Identification of the key parameters δ^{covg} and $\lambda^{ntwk}(z_i)$ comes from within-individual variation in plan characteristics over time among *plan switchers*. The primary concern for this identification strategy is that plan switching in the Connector is not exogenous; enrollees who experience a large change in unobserved health state may be most likely to switch plans (along both cost-sharing and network breadth dimensions). It is therefore a natural concern that estimates of incremental cost are likely upwards-biased. To minimize concerns due to selection into plan switching, I restrict estimation of δ^{covg} and $\lambda^{ntwk}(z_i)$ to subsamples where plan switching is predominantly driven by (plausibly) exogenous changes in relative plan prices. These subsamples are

in expected costs given observed diagnoses. I construct these risk scores using the HSS-HCC method based on current year's claims data.

[27] The ACA's metal tiers correspond to actuarial values: Bronze ˜0.6, Silver ˜0.7, Gold ˜0.8 and Platinum ˜0.9, which represents the typical share of medical costs paid by the insurer.

different for network-switchers and coverage tier-switchers, respectively, and so in practice I estimate each set of parameters separately, using Poisson regression with individual fixed effects ("xtpossion, fe" in Stata).[28]

Network switchers (estimation of $\lambda^{ntwk}(z_i)$) I estimate $\lambda^{ntwk}(z_i)$ on the unbalanced panel of patients who are continuously enrolled from the end of 2016 to the beginning of 2017, and either

1. Switch from the broad-network plan (NHP) to narrow network plans (BMC or Network Health) when the price of the broad-network plan increases in 2017, or
2. Stay enrolled in the same plan (broad or narrow network) in both years.

I further restrict to ConnectorCare Plan Type III (200-300% FPL) enrollees to ensure that that there is no concurrent variation in cost-sharing plan features among either group. The predominance of plan switching in this sample is due to the change in price: 21.8% of broad-network enrollees switch to a narrow network plan in 2017, relative to an average plan-switching rate of 1.4% among enrollees in other carriers. Identification of $\lambda^{ntwk}(z_i)$ comes from the year-on-year change in insurer costs for patients who switch from broad- to narrow-network coverage, relative to the control group of plan stayers; this is analogous to the identification of a difference-in-difference treatment effect.

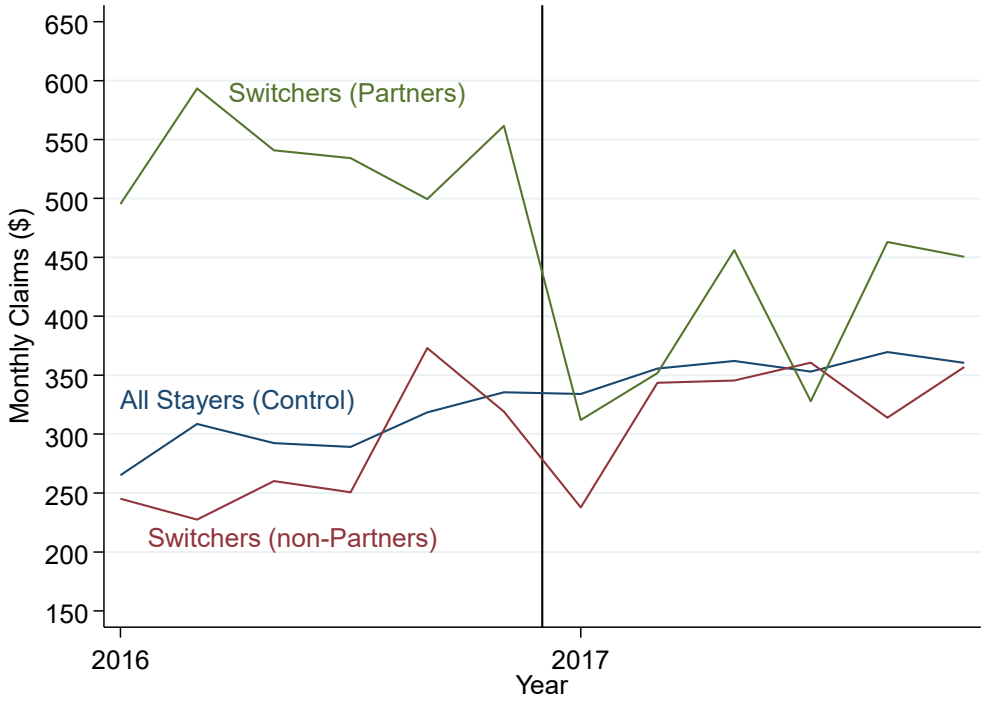
I find that broad-network coverage increases expected insurer costs by an average of 28.5% (s.e. 9.78%). Consistent with predictions that $\lambda^{ntwk}(z_i)$ is driven by patients' use of expensive and intensive providers, I find that the network cost effect is driven almost exclusively by patients who used outpatient Partners providers in 2016, prior to switching to narrow-network coverage: Partners users see a 61.4% (s.e. 18.1%) increase in expected costs when enrolled in the broad-network plan relative to a narrow-network, while non-Partners patients see a statistically insignificant 3.4% (s.e. 10.7%) increase in expected cost.[29] Partners use on the exchange is only observable for patients already enrolled in NHP; I therefore estimate heterogeneity in $\lambda^{ntwk}(z_i)$ accordig to whether each patient used outpatient Partners providers before enrolling in the exchange. I find that pre-exchange Partners patients incur a 35.8% (s.e. 15.9%) greater expected costs under broad-network coverage compared to 22.4% (s.e. 11.7%) for patient who did not use Partners prior to enrolling in the Connector. Appendix Table 3 shows the full Poisson regression results, including my final specification, which estimates the effect of NHP coverage separately by whether the patient used Partners before enrolling in the Connector, interacted with categorical measures of the patient's distance from the nearest Partners hospital.

Metal tier switchers (estimation of δ^{covg}) I focus on an unbalanced panel of broad-plan (NHP) enrollees continuously enrolled from the end of 2016 to the beginning of 2017 in order to estimate δ^{covg}. In 2017, Neighborhood Health Plan raised the price of its plans across all metal tiers, but raised

[28] Poisson regression estimators are consistent for multiplicative models such as equation 19, even if the underlying data (in this case, monthly medical spending) do not follow a Poisson distribution. I report heteroskedasticity-robust standard errors (clustered at the individual-level) to account for mis-specification of the conditional variance of medical spending. See Santos Silva and Tenreyro (2006) and for a detailed discussion.

[29] This finding also provides some reassurance that the estimates of $\lambda^{ntwk}(z_i)$ are not driven by selection on shocks to unobserved health – if that were the case, one should expect to see costs fall for non-Partners patients who switch plans.

Figure 9: Identifying Variation for Network Cost Effects



Notes: This figure shows event study estimates of $exp\left(\lambda_t^{ntwk}(z_i)\right)$ for three groups z_i: switchers who used Partners in 2016 (green line), switchers who did not use Partners in 2016 (red line) and non-switchers (blue line). The event study coefficients have also been scaled by the 2016 monthly average of insurer cost of each group.

its price *less* for its Silver-plan option relative to Bronze, Gold and Platinum. This was likely for strategic reasons, since the price of its ConnectorCare option (which accounts for a large share of its overall enrollment) is linked to the price of its Silver plan. The price of NHP Gold also fell relative to NHP Platinum. As a result, a number of NHP enrollees switch from non-Silver NHP to Silver NHP plans, and from NHP Platinum to NHP Gold: NHP enrollees switch across metal tiers at a 7.1% rate in 2017, compared to a rate of 1.1-1.5% among other carriers. I therefore construct an unbalanced panel consisting of

1. Neighborhood (NHP) enrollees who switch from Bronze, Gold or Platinum to Silver in 2017, as well as NHP enrollees who switch from Platinum to Gold, and

2. NHP enrollees who stay in the same coverage tier in 2017.

My estimate of δ^{covg} implies an increase in one coverage tier (e.g., Silver to Gold, or Gold to Platinum) increases insurer expected costs by 28.1% (s.e. 9.86%). For comparison, the ACA's risk adjustment formulas, including assumed "induced demand factors" that account for moral hazard, assume a proportional effect that ranges between 19.8-20.2% cost increases increase in metal tier. I also compute back-of-the-envelope implied arc elasticities of total claims to out-of-pocket price, approximating out of pocket price by $1 - AV$, and total claims by $\frac{1}{AV} \times \{\text{insurer cost}\}$, where AV represents the actuarial value of the metal tier. For example, the implied arc elasticity of demand of going from Platinum ($AV = 0.9$) to Silver ($AV = 0.7$) is -0.243.[30]

5 Results: Counterfactual Menu Design

5.1 Menu Design and Adverse Selection

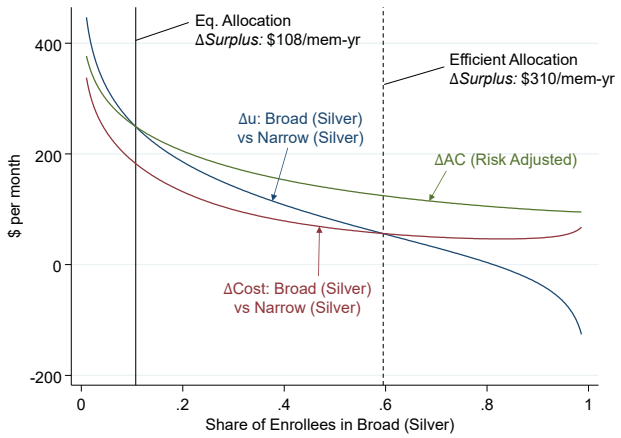
With the distribution of preferences and costs estimated in Section 4, I am equiped to empirically analyze the menu design problem described in the conceptual framework from Section 2. Following the setup in Section 2.1, I focus on two-plan menus.[31] I show results for a vertically differentiated menu, with narrow- and broad-network plans offered at the same level of financial coverage, and a diagonally differentiated menu where the broad-network plan is offered at a less generous level of financial coverage. For my main results, I use BMC as the (baseline) narrow-network plan and NHP as the broad-network plan, and exclude plan inertia from welfare analysis (assume everyone is a new enrollee). In the vertical menu, I require both plans are offered at the Silver level. For the diagonal menu, I keep Silver-tier BMC as the narrow-network plan, but make the NHP plan Bronze (less generous); Silver and Bronze are the most commonly purchased metal tiers in the Connector.

[30] The arc elasticity is given by the ratio of the percent change in total claims relative to the average, $\frac{\frac{1}{0.7}1.281 - \frac{1}{0.9}1.281^2}{\frac{1}{2}\left(\frac{1}{0.7}1.281 + \frac{1}{0.9}1.281^2\right)} \approx -0.243$ to the percent change in out-of-pocket price, $\frac{(1-0.7)-(1-0.9)}{\frac{1}{2}((1-0.7)+(1-0.9))} = \frac{0.3-0.1}{\frac{1}{2}(0.3+0.1)} = 1$.

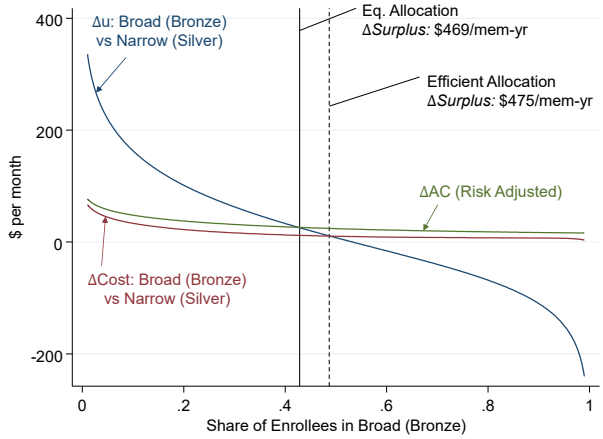
[31] Restricting to two-plan menus allows me to evaluate the efficiency of choice with straightforward graphical and computational analysis. Larger menus with potentially $N \geq 3$ plans are computationally more cumbersome without providing additional intuition.

Figure 10: Menu Design, Equilibrium, and Welfare

Panel A: Vertical Differentiation



Panel B: Diagonal Differentiation



Notes: This figure plots $\Delta Cost$ (red), Δu (blue) – the estimated (incremental) demand curve including the logit distribution term – and ΔAC (green), the risk-adjusted difference in average costs between plans, for two menus. Panel A shows a vertically differentiated menu, and Panel B shows a diagonally differentiated menu. Each Panel illustrates the efficient allocation, where incremental demand equals $\Delta Cost$, and the equilibrium allocation, where incremental demand intersects the risk-adjusted ΔAC curve. Dollar figures denoting surplus are given relative to all enrollees being enrolled in a Silver-tier broad-network plan.

Equilibrium I show market outcomes in competitive equilibrium following the notion of price competition of Handel, Hendel and Whinston (2015) between insurers offering one of two pre-determined insurance contracts. Insurers earn zero profits in equilibrium, meaning the equilibrium is characterized by the break-even price, where the relative price between plans equals the difference in average insurer costs.

Risk Adjustment In computing equilibrium prices, I include a realistic form of risk adjustment, which is similar to that used in the ACA. In both vertical and diagonal menus, plans receive risk-adjustment payments for each their enrolled patients which are equal to that patient's expected costs (based on concurrent risk scores) *using the narrow-network plan as a baseline.* This risk adjustment scheme essentially leads the broad-network insurer to internalize the savings of the narrow-network that result from changes in the relative price of the broad-network plan.

Results Figure 10 shows the empirical analog to Figure 2 in Section 2.3. In each panel, market equilibrium is given by where the blue (demand) curve intersects the green (risk-adjusted ΔAC curve), while the socially efficient price is given by where the demand curve intersects the red (marginal cost) curve.Panel A on the left shows the result for a vertically differentiated menu, when both narrow- and broad-network plans are offered at the same Silver-tier coverage level. The marginal cost curve (illustrated in red) is steep, which generates severe adverse selection. Despite risk adjustment, the broad-network plan unravels substantially. Panel B on the right shows the result for the diagonally differentiated menu. The key effect of diagonal differentiation is to flatten the marginal cost curve. This increases market surplus at the efficient allocation, since it increases the area between the demand and

Table 1: Equilibrium Market Efficiency Under Various Menus

	Two-Plan Menus		Characteristics of Equilibrium Allocation (1)	(2)	(3)	(4)
	Narrow (Tier)	vs Broad (Tier)	Broad Plan Share	Rel. Price Broad ($/yr)	Rel. Surplus ($/member-yr) excl. Logit	Rel. Surplus ($/member-yr) incl. Logit
i	Narrow (Bronze)	Broad (Bronze)	0.328	1572	-103.0	138.0
ii	Narrow (Silver)	Broad (Silver)	0.107	2988	-27.9	107.5
iii	Narrow (Bronze)	Broad (Silver)	0.003	8208	-264.3	-255.2
iv	Narrow (Silver)	Broad (Bronze)	0.427	312	190.3	469.4

Notes: Table describes characteristics of the equilibrium allocation (with ACA-like risk adjustment) within two-plan menus described in each row. Columns 1 and 2 describe the allocation and equilibrium price (in $/year), respectively. Column 3 gives the relative social surplus of the optimal allocation, compared to a Silver-tier broad-network plan, when the sponsor's objective is evaluated excluding the logit distribution term. Column 4 evaluates the same allocation relative to a Silver-tier broad-network plan under a sponsor's objective that includes the logit term.

marginal cost curves. It also leads to substantially less unraveling than in the vertical menu. In this sense, one can think of the diagonal differentiation as acting as a built-in form of risk adjustment, by compensating the broad-network provider indirectly (by allowing it to cover a smaller share of costs) rather than directly (through more sophisticated risk adjustment). The idea of combating adverse selection through menu design in this way dates back at least to Enthoven (1988).

Table 1 shows detailed characteristics of market equilibria for a variety of two-plan menus, including the vertical and diagonal menus shown above in Figure 10. The diagonal menu (row iv) achieves an equilibrium allocation nearly as efficient as the optimal allocation, generating close to 100 percent of the surplus of the optimal allocation. However, the vertical menu (row ii) generates a much smaller surplus in equilibrium than at the efficient allocation, due to adverse selection, generating only $\frac{107.5}{310.8} \approx 34.6$ percent of the surplus at the efficient allocation.

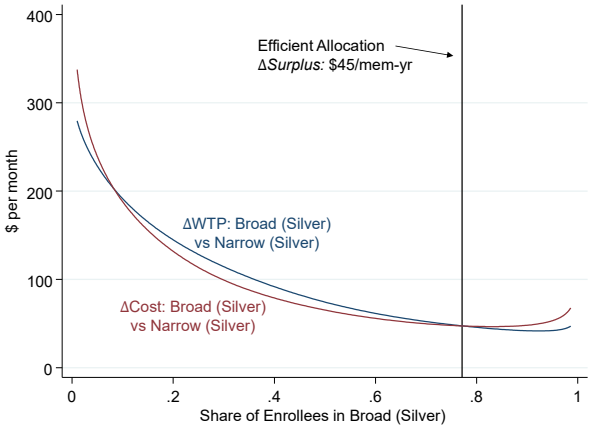
5.2 Additional Analyses and Robustness

Interpretation of Logit Term ϵ_{ijt} The unobserved logit error term is typically interpreted as capturing unobserved preference heterogeneity. Because the specification assumes draws of the unobserved preference term are i.i.d. across plans j for each choice instance, the logit specification is notorious for implying "new product" welfare effects which may overstate the value of expanding the choice set. This provides additional motivation to restricting analysis to two-plan menus, where the size of choice sets is held constant across counterfactual menus. Another challenge for interpretation of the logit distribution, which is specific to the ACA marketplace setting, is that there may be unobserved differences in the *price* faced by consumers due to external premium subsidies, e.g. employer reimbursements.[32] Variance in unobserved premium reimbursements would tend to make interpreting the logit distribution as unobserved preferences overstate the extent of preference heterogeneity. As a robustness exercise, I report welfare results that entirely exclude the logit term from the welfare

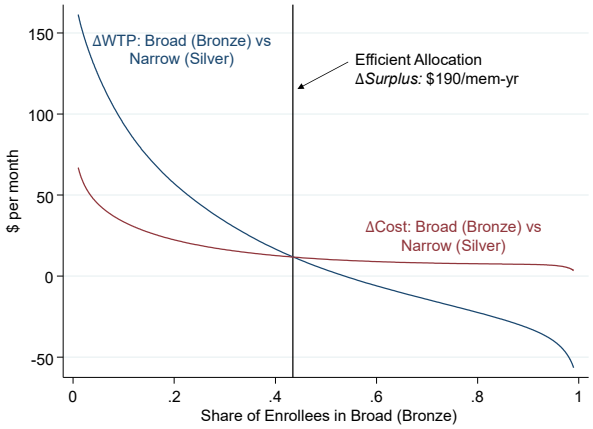
[32]Such reimbursements were initially prohibited under the ACA, but became legal for small employers in 2017.

Figure 11: Robustness: Excluding Logit Term from Welfare

Panel A: Vertical Differentiation



Panel B: Diagonal Differentiation



Notes: The figures plot the average $\Delta Cost$ (red) and ΔWTP (blue) of marginal enrollees against the allocations generated by varying the relative price for vertically and diagonally differentiated menus. ΔWTP excludes the logit term from the sponsor's objective.

analysis. In this exercise, I assume welfare includes only the predicted (average) WTP from Equation (18) to evaluate welfare.[33] I consider these approaches as giving lower bounds on the value of choice.[34]

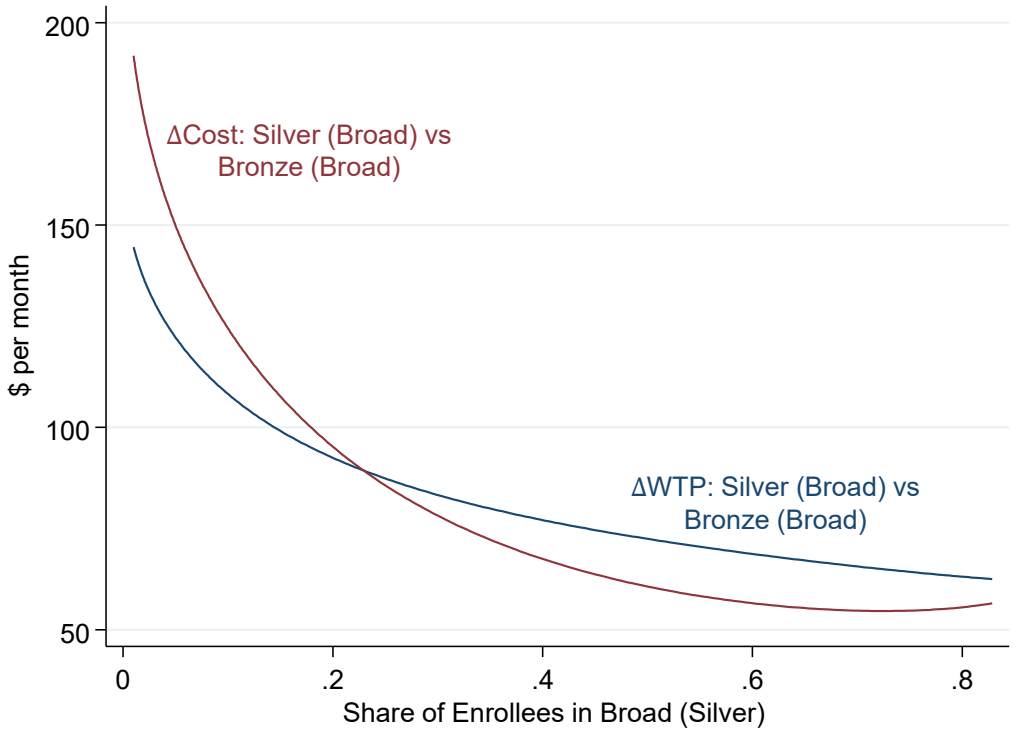
Panel A of Figure 11 plots $\Delta Cost$ and ΔWTP for vertical differentiation in network quality. Again, differentiation in network quality induces selection on incremental costs. However, in this case, there is an interior allocation that maximizes the sponsor's objective while offering choice over plans with different networks: this can be achieved with a relative price of \$12 per month (\$144 per year), which induces 77.1 percent of enrollees to take up the Silver-Broad plan.Panel B plots $\Delta Cost$ and ΔWTP curves for a diagonally differentiated menu consisting of a Bronze-tier broad-network plan and a Silver-tier narrow-network plan. The key qualitative difference, compared with the examples of vertical differentiation shown above, is that this menu generates less selection on incremental cost, which can be seen from the relatively shallow slope of the $\Delta Cost$ curve. At the same time, this menu offers variety that appeals differently to different patient types. The planner's best allocation offers choice between these two plans with a relative price of \$24 per month (\$288 per year) resulting in 43.5 percent of enrollees choosing the Bronze-Broad plan.

Vertical Differentiation in Financial Coverage Figure 12 shows incremental WTP (excluding the Logit term) and marginal cost curves for the population of patients on the margin at each allocation. Several things are apparent from the figure. The steep downward slope of the marginal cost ($\Delta Cost$) curve indicates high selection on (incremental) marginal cost. In fact, at higher prices, Silver is taken

[33]Notably, the predicted WTP excludes preference heterogeneity attributed to predictable plan inertia. Inertia plays an important role in identifying preferences for plan characteristics, but interpreting it as preference heterogeneity runs into the well-known problem of distinguishing preference heterogeneity from state dependence. I therefore remove the plan inertia term from my welfare calculations, in effect treating all enrollees as if they are first-time enrollees.

[34]Non-welfare-relevant logit errors effectively act as noise in the assignment of plans; in the extreme case where plan assignment is entirely uncorrelated with predicted WTP, offering choice would be weakly worse than only offering the plan that is most efficient on average.

Figure 12: Vertical Differentiation in Financial Coverage



Notes: This figure plots $\Delta Cost$ and ΔWTP (excluding logit term in sponsor's objective) for marginal enrollees against the allocations generated by varying the relative price of plans that are vertically differentiated in financial coverage. ΔWTP excludes the logit term from demand.

Table 2: Efficient Allocations in Two-Plan Menus

	Two-Plan Menus		Characteristics of Efficient Allocation (1) Offers Choice?	(2) Broad Plan Share	(3) Rel. Price Broad ($/yr)	(4) Rel. Surplus ($/member-yr)
	Narrow (Tier) vs	Broad (Tier)				
A. Planner's Objective Excludes Logit						
i	Narrow (Bronze)	Broad (Bronze)	No	1	N/A	-180.0
ii	Narrow (Silver)	Broad (Silver)	Yes	0.771	144	45.0
iii	Narrow (Bronze)	Broad (Silver)	No	1	N/A	0
iv	Narrow (Silver)	Broad (Bronze)	Yes	0.435	288	190.3
B. Planner's Objective Includes Logit						
i	Narrow (Bronze)	Broad (Bronze)	Yes	0.661	480	257.1
ii	Narrow (Silver)	Broad (Silver)	Yes	0.595	672	310.8
iii	Narrow (Bronze)	Broad (Silver)	Yes	0.811	756	156.1
iv	Narrow (Silver)	Broad (Bronze)	Yes	0.487	132	474.6

Notes: Table describes characteristics of the optimal allocation within two-plan menus described in each row. Columns 1-3 describe the efficient allocation: whether it offers choice (Column 1), what share of enrollees are assigned the Broad-network plan (Column 2), and the Δp for the Broad-network plan which achieves the allocation. Column 4 gives the relative social surplus of the optimal allocation, evluated relative to a single-plan menu consisting of only a Silver-tier Broad-network plan. Panel A characterizes the efficient allocation when the sponsor's objective excludes the logit error term, while Panel B characterizes the efficient allocation in the same set of menus under an objective that includes the logit error term.

up by patients for whom the relative cost of providing them with Silver on average exceeds their relative willingness-to-pay for Silver, which is shown graphically by the $\Delta Cost$ curve crossing the ΔWTP curve from above. As a result, there is no price at which the market sponsor efficiently offers choice between Bronze and Silver plans; the sponsor's best allocation is achieved by restricting enrollees to Silver.

Optimal menu I evaluate the socially optimal allocation under all possible two-plan menus. Table 2 gives the results. Panel A characterizes the efficient allocations when the sponsor's objective excludes the logit term. The diagonally differentiated menu described above (Bronze-Broad and Silver-Narrow) is the optimal two-plan menu. When optimally price, it generates approximately $475 per member-year ($190 excluding the logit term) greater social surplus than a single-plan menu consisting of only a Silver-Broad plan, and roughly $165 per member-year ($145.3 excluding the logit term) more than the best vertically differentiated menu, which offers choice between Narrow and Broad Silver plans.

6 Conclusion

Adverse selection poses a fundamental threat to insurance markets. It can distort the incentives of both insurance sellers and buyers from making socially efficient choices. To address adverse selection, economists have studied a variety of corrective policy tools, principally *financial incentives* like risk adjustment and subsidies or mandates.

This paper highlights another set of policy tools: *menu design*. Menu design policies govern the

kinds of plans that are allowed to be traded, and are widely employed in many health insurance settings. Notably, these include the ACA exchanges and Medicare Advantage, as well as numerous employer-sponsored health insurance menus.

I argue that menu design is an important policy margin which can be used, in addition to financial incentives, to address adverse selection. Menu design is important because the severity of adverse selection depends on the *kinds of choices* available in the market. I formally illustrate this argument with a stylized model of an insurance market, and also show the idea graphically in the selection markets framework of (Einav, Finkelstein and Cullen, 2010). I then leverage policy-driven price variation in an important policy setting, the Massachusetts ACA exchange, to estimate a model of insurance demand and cost which is closely related to the conceptual model. The analysis finds that requiring plans to differentiate in only a single quality dimension can exacerbate adverse selection. However, offering *diagonally differentiated* choice – offsetting two quality dimensions – may substantially alleviate selection problems.

Menu design has both demand-side and supply-side effects on adverse selection. On the demand-side, it can be thought of as an approach to the multi-dimensional screening problem which arises because prices cannot be used to screen patients into socially-efficient coverage. Thus, menu design is relevant in non-market insurance settings when plans are provided and prices can be determined by a central regulator or planner, such as in some employer-designed menus. The paper also highlights menu design's role as a supply-side policy in competitive insurance markets, where it can be thought of as acting similarly to a form of flexible and automatic risk adjustment to address market unraveling. In my empirical analysis, I find that a diagonally differentiated menu leads to both more efficient consumer sorting and less severe menu unraveling, but it is important to note that this may not be the case in other settings. The conceptual framework provides guidance to understanding potential *tradeoffs* between demand- and supply-side effects of different menus. The framework can also be used to analyze policies other than menu design which affect how consumers sort between plans. For example, Handel (2013) studies a setting in reducing plan inertia involves a tradeoff between consumer choice efficiency and supply-side adverse selection.

This paper may help explain several facts about real-world health insurance markets. Restrictions on differentiation in financial coverage are a possible driver of adverse selection leading to the prevalence of narrow-network coverage in ACA exchanges. The findings also have implications for how policymakers think about regulation in insurance markets. Policies that restrict how insurance plans can differentiate from each other may have unintended market consequences because of their effects on adverse selection. Quantifying tradeoffs between menu design's effects on adverse selection versus other possible benefits, such as simplifying choices for consumers, is an interesting area for future work.

References

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

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

Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–3248.

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

Cutler, David M. and Sarah J. Reber, "Paying For Health Insurance: The Trade-off Between Competition and Adverse Selection," *The Quarterly Journal of Economics*, May 1998.

Dafny, Leemore, Igal Hendel, and Nathan Wilson, "Narrow Networks on the Health Insurance Exchanges: What Do They Look Like and How Do They Affect Pricing? A Case Study of Texas," *American Economic Review: Papers and Proceedings*, 2015, *105* (5), 110–114.

_ , Kate Ho, and Mouricio Varela, "Let Them Have Choice: Gains from Shifting Away from Employer-Sponsored Health Insurance and Toward an Individual Exchange," *American Economic Journal: Economic Policy*, 2013, *5* (1), 32–58.

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

_ , _ , and Neale Mahoney, "The IO of Selection Markets," in Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization*, Elsevier, 2021.

_ , _ , Stephen P. Ryan, Paul Schrimpf, and Mark Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, *103* (1), 178–219.

Enthoven, Alain C., *Theory and Practice of Managed Competition in Health Care Finance*, Elsevier, 1988.

Ericson, Keith Marzilli and Amanda Starc, "Measuring Consumer Valuation of Limited Provider Networks," *American Economic Review*, 2015, *105* (5), 115–119.

_ and _ , "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," *Review of Economics and Statistics*, 2015, *97* (3), 667–682.

_ and _ , "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics*, 2016, *50*, 71–85.

Finkelstein, Amy and James Poterba, "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market," *Journal of Political Economy*, 2004, *112* (1).

_ **and Kathleen McGarry**, "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, September 2006, *96*, 938–958.

Geruso, Michael and Timothy Layton, "Selection in Health Insurance Markets and Its Policy Remedies," *Journal of Economic Perspectives*, 2017, *31* (4), 23–50.

_ **,** _ **, and Daniel Prinz**, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," *American Economic Journal: Economic Policy*, May 2019, *11* (2).

_ **,** _ **, Grace McCormack, and Mark Shepard**, "Two Margin Problem in Insurance Markets," *Review of Economics and Statistics*, 2021, *Forthcoming.*

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

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

Handel, Benjamin, "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682.

Ho, Kate and Robin Lee, "Health Insurance Menu Design for Large Employers," *The RAND Journal of Economics*, 2022.

Jaffe, Sonia and Mark Shepard, "Price-Linked Subsidies and Imperfect Competition in Health Insurance," *American Economic Journal: Economic Policy*, 2020, *12* (3), 279–311.

Landais, Camille, Arash Nekoei, Peter Nilsson, David Seim, and Johannes Spinnewijn, "Risk-Based Selection in Unemployment Insurance: Evidence and Implications," *American Economic Review*, 2021, *111* (4), 1315–1355.

_ **, Nathaniel Hendren, and Johannes Spinnewijn**, "Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design," *Annual Review of Economics*, 2022.

Lavetti, Kurt and Kosali Simon, "Strategic Formulary Design in Medicare Part D Plans," *American Economic Journal: Economic Policy*, August 2018, *10* (3).

Marone, Victoria and Adrienne Sabety, "When Should There Be Vertical Choice in Health Insurance Markets?," *American Economic Review*, 2022, *112* (1), 304–342.

Nevo, Aviv, "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," *RAND Journal of Economics*, 2000, *31* (3), 395–421.

Panhans, Matthew, "Adverse Selection in ACA Exchange Markets: Evidence from Colorado," *American Economic Journal: Applied Economics*, 2019, *11* (2).

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

Santos Silva, J.M.C. and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, November 2006, *88* (4), 641–658.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," *American Economic Review*, 2022, *112* (2), 578–615.

Tilipman, Nicholas, "Employer Incentives and Distortions in Health Insurance Design: Implications for Welfare and Costs," *American Economic Review*, 2022, *112* (3), 998–1037.

A Theory Appendix

A.1 Welfare and Choice Efficiency

Begin with the definition of the gains from choice in Equation (8). For a given price Δp, and suppressing the dependence of s and $\overline{\gamma}_1$ on Δp, note that

$$\begin{aligned}
s \cdot \overline{\gamma}_1 &= s \cdot \overline{\gamma}_1 + \overline{\gamma} - \overline{\gamma} \\
&= s \cdot \overline{\gamma}_1 + [s \cdot \overline{\gamma}_1 + (1-s) \cdot \overline{\gamma}_0] - \overline{\gamma} \\
&= s\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + s \cdot \overline{\gamma}_1 + \overline{\gamma}_0 - [s \cdot \overline{\gamma}_1 + (1-s) \cdot \overline{\gamma}_0] \\
&= s(1-s)\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + s^2\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + \overline{\gamma}_0 - (1-s) \cdot \overline{\gamma}_0 \\
&= s(1-s)\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + s^2\overline{\gamma}_1 + s(1-s)\,\overline{\gamma}_0 \\
&= s(1-s)\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + s[s \cdot \overline{\gamma}_1 + (1-s) \cdot \overline{\gamma}_0] \\
&= s(1-s)\left(\overline{\gamma}_1 - \overline{\gamma}_0\right) + s \cdot \overline{\gamma}.
\end{aligned}$$

Let $\delta_i = 1\left(\Delta V_i \geq \Delta p\right)$ be an indicator for whether i buys plan H and price Δp, and recall the definition of $\overline{\gamma}_1\left(\Delta p\right) = E\left[\gamma_i \,|\, \Delta V_i \geq \Delta p\right] = E\left[\gamma_i \,|\, \delta_i = 1\right]$, and $\overline{\gamma}_0\left(\Delta p\right) = E\left[\gamma_i \,|\, \Delta V_i < \Delta p\right] = E\left[\gamma_i \,|\, \delta_i = 0\right]$. Since δ is a binary indicator variable for selecting H, the difference in conditional means $\overline{\gamma}_1 - \overline{\gamma}_0$ is equivalent to the slope coefficient from a linear regression of γ on δ, i.e.,

$$\overline{\gamma}_1 - \overline{\gamma}_0 = \frac{\mathrm{Cov}\left(\gamma, \delta\right)}{\sigma_\delta^2}.$$

Note also that, since δ is a Bernoulli random variable which equals 1 with probability $s\left(\Delta p\right)$, its variance $\sigma_\delta^2 = s\left(\Delta p\right) \cdot \left(1 - s\left(\Delta p\right)\right)$. We therefore have

$$s \cdot \overline{\gamma}_1 = \mathrm{Cov}\left(\gamma, \delta\right) + s \cdot \overline{\gamma}. \tag{20}$$

The next step is to show the relationship between $\mathrm{Cov}\left(\gamma, \delta\right)$ and choice efficiency, $\mathrm{Cov}\left(\gamma, \Delta V\right)$. Consider the linear projection of γ onto ΔV:

$$\gamma_i = \frac{\mathrm{Cov}\left(\gamma, \Delta V\right)}{\sigma_{\Delta V}^2} \cdot \Delta V + \tilde{\gamma}_i, \tag{21}$$

where $\tilde{\gamma}_i$ is an idiosyncratic displacement term.

Plugging Equation (21) into $\mathrm{Cov}\left(\gamma, \delta\right)$ yields

$$\begin{aligned}
\mathrm{Cov}\left(\gamma, \delta\right) &= \mathrm{Cov}\left(\tfrac{\mathrm{Cov}(\gamma, \Delta V)}{\sigma_{\Delta V}^2} \cdot \Delta V + \tilde{\gamma}\,,\, \delta\right) \\
&= \tfrac{\mathrm{Cov}(\gamma, \Delta V)}{\sigma_{\Delta V}^2} \cdot \mathrm{Cov}\left(\Delta V\,,\, \delta\right) + \mathrm{Cov}\left(\tilde{\gamma}, \delta\right).
\end{aligned}$$

Note that, for any random variable X,

$$\mathrm{Cov}\left(X, \delta\right) = s\left(1-s\right) \cdot \left\{E\left[X \,|\, \delta = 1\right] - E\left[X \,|\, \delta = 0\right]\right\},$$

since δ is an indicator variable. So we are left with

$$\text{Cov}(\gamma, \delta) = \delta V(\Delta p) \cdot \text{Cov}(\gamma, \Delta V) + \delta\tilde{\gamma}(\Delta p), \tag{22}$$

where $\delta V(\Delta p) = \frac{s(1-s)}{\sigma^2_{\Delta V}} \cdot \{E[\Delta V \mid \delta = 1] - E[\Delta V \mid \delta = 0]\}$, and $\delta\tilde{\gamma}(\Delta p) = s(1-s) \cdot \{E[\tilde{\gamma} \mid \delta = 1] - E[\tilde{\gamma} \mid \delta = 0]\}$

The final step is to plug Equations (20) and (22) into Equation (8), which yields,

$$SW(\Delta p) = SW_L \ + \ s(\Delta p) \cdot \overline{\gamma} \ + \ \delta V(\Delta p) \cdot \text{Cov}(\gamma, \Delta V) \ + \ \delta\tilde{\gamma}(\Delta p).$$

B APCD Details and Data Construction

B.1 Enrollment and Claims Data

The APCD allows me to observe or infer a number of key variables, including:

- **Enrollee Demographics:** the APCD includes variables describing each individual's household membership, age, sex, and location (zip code). Based on each enrollee's zip code, I construct measures of the average distance to a number of key Partners hospitals.

- **Plan Characteristics:** based on an enrollee's household income, he/she may be eligible for premium and cost-sharing subsidies through the ConnectorCare program. The APCD does not include information on individual- or household-level income, but does include information on plan characteristics including each plan's actuarial value. Based on the exchanges income-based subsidy schedule, I infer subsidy eligibility from the actuarial value of the each enrollee's plan.

- **Individual Health Risk:** sickness is likely a key driver of plan selection. I construct a measure of each enrollee's health ("risk score") using medical diagnoses listed on medical claims in the APCD, following the HHS-HCC risk score definition. The HCC risk score is a *concurrent* measure of risk, based on each patient's current-year medical diagnoses. While I observe plan enrollment for continuing enrollees in 2018, the APCD claims I have include only 2013-2017. For plan choices made in 2018, I use each patient's 2017 risk score as a proxy for health at the time of choice.

- **Provider Usage History:** an advantage of the APCD is that it allows me to observe both enrollment and claims data for individuals *outside* of the exchange. I use claims for enrollees before they first enroll in the exchange to create a measure of loyalty to Partners Healthcare providers, based on whether each patient used outpatient Partners up to 2 years before entering the exchange.[35] Unlike any measure of Partners use after enrolling in the exchange, this is not endogenous to plan choice on the Connector. I identify Partners-affiliated physicians and outpatient facilities using the Massachusetts Provider Database from Massachusetts Health Quality Partners (MHQP).

[35] Since I observe claims as early as 2013, and the exchange first opens in 2015, two years is the greatest window within which I can observe pre-exchange Partners use for all enrollees.

B.2 Construction of Plan Choice Dataset

I make a number of sample restrictions in the final plan choice dataset.

1. Restrict sample to Connector QHP individual (household size of one) enrollees not eligible for ConnectorCare plans (>300% of FPL), for which I directly observe relative plan prices.

2. Restrict choice set to Bronze, Silver, Gold and Platinum plans offered by Neighborhood Health Plan, BMC Healthnet, and Network Health. These three carriers collectively make up the large majority of enrollment in ConnectorCare, but there are two major carriers (Harvard Pilgrim and Blue Cross Blue Shield) that do not participate in ConnectorCare but do offer QHPs in the individual market. I exclude these carriers because while they are broad network plans, I do not directly observe their hospital networks, and because there is almost no margin of substitution between them and the three main carriers I study.

3. Restrict sample to enrollees who live in Massachusetts Connector rating areas 3-6. The Connector allows carriers to vary the price of their plans according to an individual's age and the "rating area" in which they live. Massachusetts is divided into seven rating areas, characterized by the first 3 digits of the location's zip code. I exclude rating areas 1 (3-digit zips 010, 011, 012 and 013) and 2 (014, 015 and 016) which capture Western Massachusetts, as well as rating area 7 (3-digit zip codes 025 and 026) which covers Cape Cod and Martha's Vineyard/Nantucket. The reason for excluding these areas is that BMC's provider network has very sparse coverage in these regions and, while I do observe non-zero shares of some BMC plans among enrollees in those zip codes, it is unclear whether those choices should inform demand estimates.

4. Collapse plan choices to one option per metal tier per carrier, for a total of up to 12 available choices. While the Connector requires each carrier to offer a standardized cost-sharing plan in each metal tier, in practice each carrier in my sample offers more than one plan in some metal tiers in some years. These additional plans are allowed to be "non-standard," with slightly different cost-sharing and prescription drug formularies than the exchange's standard requirements, although they are subject to the regulatory approval of the exchange. Unfortunately, I do not observe the contract characteristics of the non-standard plans, and so cannot control for differences in contract features across different plans within carrier-metal tier. Instead, I collapse all choices within carrier-metal tier-year into a single choice, and assign it the price of the cheapest such plan. Likely as a result of regulatory oversight restricting the carriers' abilities to differentiate non-standard plans, the prices of all plans within carrier-metal tier-year do not differ widely.

C Identification Strategy

C.1 Additive Separability of Plan Utility

This section discusses the additive separability assumption underlying the identification strategy described in Section 4.1. One way to found additive separability is to begin with a two-period model as in Einav et al. (2013). In the second period, patients have already been enrolled in a plan j, characterized by generosity and network (m_j, x_j). They receive a health shock and make a healthcare utilization decision that takes into account the features of their health plan: the plan's financial generosity and provider network. In particular, I model care as consisting of the *quantity* of care (consisting of a number and type of visits) q, and the *identity* of the healthcare providers ϕ. At the beginning of the coverage period (one year in the case of the MA Connector) patients draw a health shock λ and a provider preference shock γ from distributions known by the individual.

I assume the second-period utility function over health is additively separable in the quantity of care and the identity of providers, so that these decisions are made independently: patients trade off the value of healthcare against the value of non-healthcare consumption in order to determine the quantity of care they consume, and choose providers in order to maximize their utility from provider identity. This is a natural assumption if, as is the case in the Connector, plans of the same coverage level m all have the same out-of-pocket price schedule. Formally, second period utility is given by

$$u\left(q, \phi; \lambda, \gamma, j\right) = v\left(q; \lambda, m_j\right) + h\left(\phi; \gamma, x_j\right),$$

where utility in healthcare quantity depends on both the health shock λ and the plan generosity m_j, and utility in provider identity depends on the provider preference shock γ and the plan's provider network x_j. It thus follows that plan utility, given by the expectation (over the joint distribution of health and provider preference shocks) of second-period utility, is also additively separable in the plan characteristics m and x.

Figure 13: Connector Standardized Metal Tier Designs, 2017

Plan Feature/ Service *A check mark (✓) indicates that this benefit is subject to the annual deductible*		Platinum	Gold	Silver	Bronze
Annual Deductible – Combined		N/A	$1,000	$2,000	N/A
		N/A	$2,000	$4,000	N/A
Annual Deductible – Medical		N/A	N/A	N/A	$2,750
		N/A	N/A	N/A	$5,500
Annual Deductible – Prescription Drugs		N/A	N/A	N/A	$250
		N/A	N/A	N/A	$500
Annual Out-of-Pocket Maximum		$3,000	$5,000	$7,150	$7,150
		$6,000	$10,000	$14,300	$14,300
Primary Care Provider (PCP) Office Visits		$25	$30	$30	$25 ✓
Specialist Office Visits		$40	$45	$50	$40 ✓
Emergency Room		$150	$150 ✓	$700 ✓	$500 ✓
Urgent Care		$40	$45	$50	$40 ✓
Inpatient Hospitalization		$500	$500 ✓	$1,000 ✓	$1,000 ✓
Skilled Nursing Facility		$500	$500 ✓	$1,000 ✓	$1,000 ✓
Durable Medical Equipment		20%	20% ✓	20% ✓	20% ✓
Rehabilitative Occupational and Rehabilitative Physical Therapy		$40	$45	$50	$40 ✓
Laboratory Outpatient and Professional Services		$0	$20 ✓	$25 ✓	$50 ✓
X-rays and Diagnostic Imaging		$0	$20 ✓	$25 ✓	$175 ✓
High-Cost Imaging		$150	$200 ✓	$500 ✓	$1,000 ✓
Outpatient Surgery: Ambulatory Surgery Center		$500	$250 ✓	$750 ✓	$750 ✓
Outpatient Surgery: Physician/Surgical Services		$0	$0 ✓	$0 ✓	$0 ✓
Prescription Drug	Retail Tier 1	$15	$20	$20	$25 ✓
	Retail Tier 2	$30	$30	$60	$75 ✓
	Retail Tier 3	$50	$50	$90	$100 ✓
	Mail Tier 1	$30	$40	$40	$50 ✓
	Mail Tier 2	$60	$60	$120	$150 ✓
	Mail Tier 3	$150	$150	$270	$300 ✓
2017 Final FAVC		91.73%	81.43%	71.84%	61.86%

31

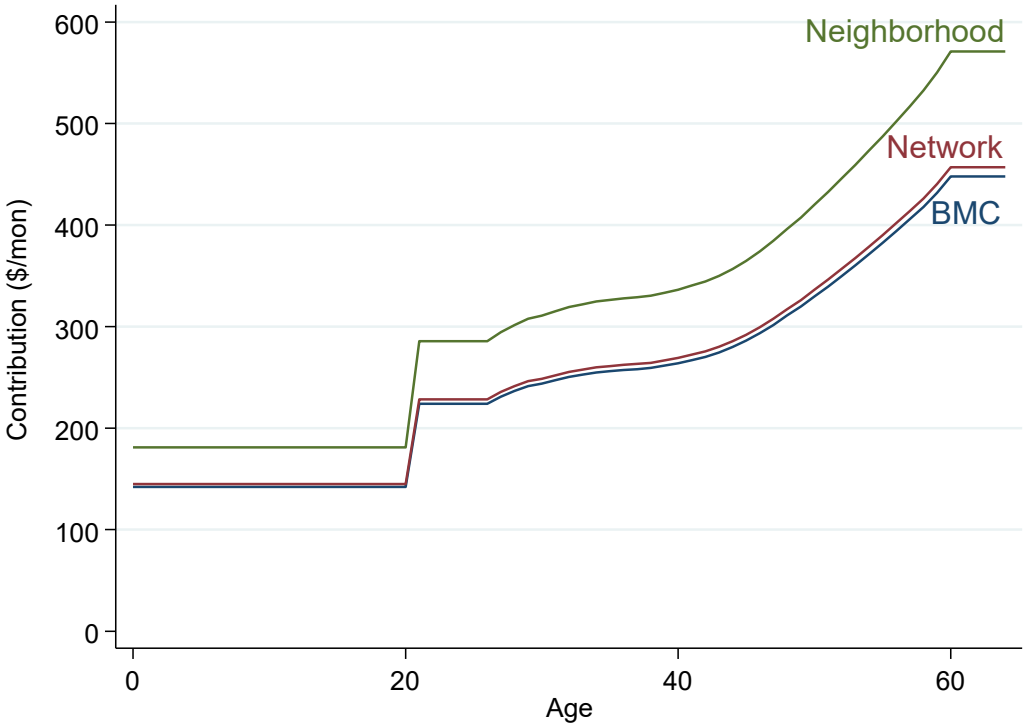
Notes: From MA Connector Final Award of 2017 Seal of Approval meeting slides.

Figure 14: Connector Standardized Metal Tier Designs, 2018

Plan Feature/ Service *A check mark (✓) indicates that this benefit is subject to the annual deductible*		Platinum	Gold	Silver	Bronze #1	Bronze #2 (HSA)
Annual Deductible – Combined		$0	N/A	$2,000	$2,500	$3,000
		$0	N/A	$4,000	$5,000	$6,000
Annual Deductible – Medical		N/A	$1,000	N/A	N/A	N/A
		N/A	$2,000	N/A	N/A	N/A
Annual Deductible – Prescription Drugs		N/A	$0	N/A	N/A	N/A
		N/A	$0	N/A	N/A	N/A
Annual Out-of-Pocket Maximum		$3,000	$5,000	$7,350	$7,350	$6,650
		$6,000	$10,000	$14,700	$14,700	$13,300
Primary Care Provider (PCP) Office Visits		$20	$30	$30	$30 ✓	$20 ✓
Specialist Office Visits		$40	$45	$50	$50 ✓	$40 ✓
Emergency Room		$150	$150 ✓	$700 ✓	$700 ✓	$250 ✓
Urgent Care		$40	$45	$50	$50 ✓	$40 ✓
Inpatient Hospitalization		$500	$500 ✓	$1,000 ✓	$1,000 ✓	$750 ✓
Skilled Nursing Facility		$500	$500 ✓	$1,000 ✓	$1,000 ✓	$750 ✓
Durable Medical Equipment		20%	20% ✓	20% ✓	20% ✓	20% ✓
Rehabilitative Occupational and Rehabilitative Physical Therapy		$40	$45	$50	$50 ✓	$40 ✓
Laboratory Outpatient and Professional Services		$0	$20 ✓	$25 ✓	$25 ✓	$25 ✓
X-rays and Diagnostic Imaging		$0	$20 ✓	$25 ✓	$25 ✓	$25 ✓
High-Cost Imaging		$150	$200 ✓	$500 ✓	$500 ✓	$500 ✓
Outpatient Surgery: Ambulatory Surgery Center		$250	$250 ✓	$750 ✓	$750 ✓	$500 ✓
Outpatient Surgery: Physician/Surgical Services		$0	$0 ✓	$0 ✓	$0 ✓	$0 ✓
Prescription Drug	Retail Tier 1	$10	$20	$20	$20	$20 ✓
	Retail Tier 2	$25	$30	$60	$60 ✓	$40 ✓
	Retail Tier 3	$50	$50	$90 ✓	$90 ✓	$60 ✓
	Mail Tier 1	$20	$40	$40	$40	$40 ✓
	Mail Tier 2	$50	$60	$120	$120 ✓	$80 ✓
	Mail Tier 3	$150	$150	$270 ✓	$270 ✓	$180 ✓
2018 Final FAVC		88.24%	79.69%	71.40%	64.84%	64.88%

Notes: From MA Connector Final Award of 2018 Seal of Approval meeting slides. New in 2018, the Connector provided an option for carriers to offer a second standardized Bronze plan option that was HSA compatible (Bronze #2). None of Neighborhood (NHP), BMC Healthnet or Network Health offered a Bronze #2 plan design in 2018.

Figure 15: Massachusetts Connector Pricing Function



Notes: Figure plots the ACA plan pricing function by age ("age curve"), illustrated by each carrier's Silver-tier plan offered in 2016 in Cambridge, MA (rating area 5). Insurance carriers have discretion over the level of price for each plan they offer within rating area × year, but the shape of the pricing schedule as a function of age is fixed across plans (and set by state regulation).

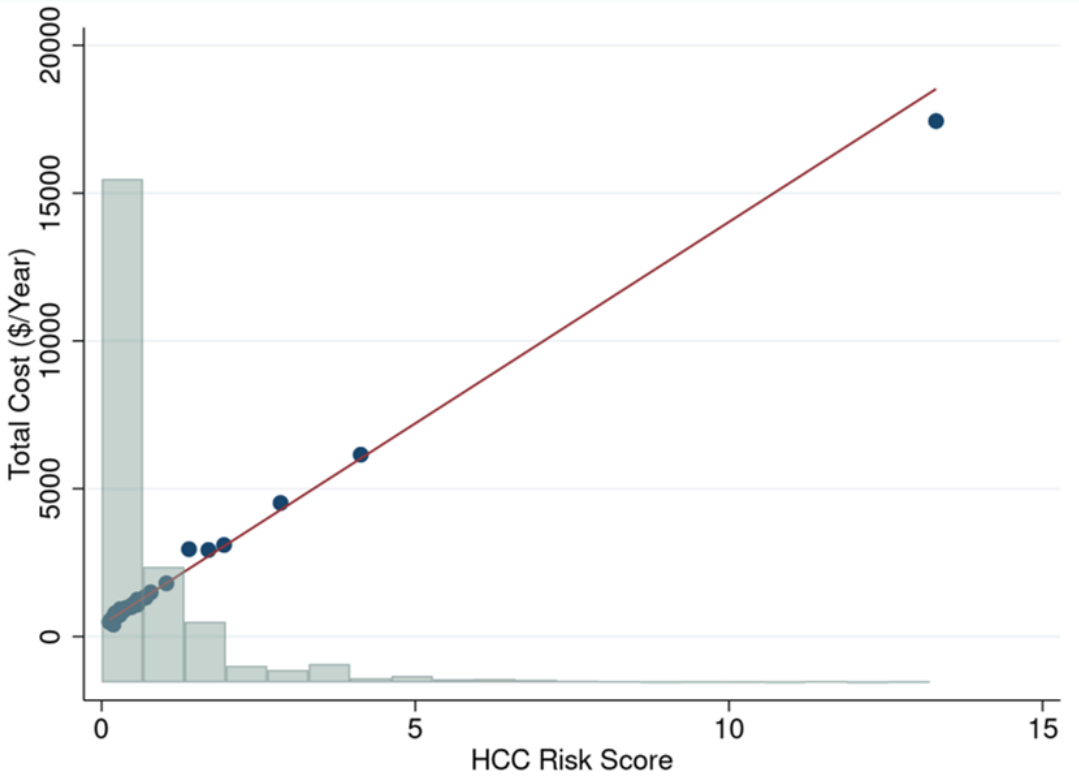
D Cost Effect Estimates

Table 3: Network Cost Effect Estimates

	Network Cost Specification			
Coefficient	(1)	(2)	(3)	(4)
Proportional Effect of NHP	1.285			
	(0.098)			
x Non-Partners (pre-Switch)		1.034		
		(0.107)		
x Partners (pre-Switch)		1.613		
		(0.181)		
x Non-Partners (pre-Connector)			1.224	
			(0.117)	
x Distance > 25 mi.				1.212
				(0.183)
x Distance 5-25mi.				1.161
				(0.144)
x Distance <5 mi.				1.605
				(0.468)
x Partners (pre-Connector)			1.358	
			(0.159)	
x Distance > 25 mi.				1.307
				(0.358)
x Distance 5-25mi.				1.342
				(0.173)
x Distance <5 mi.				1.521
				(0.262)
Time FEs	Yes	Yes	Yes	Yes
N	47,460	47,460	47,460	47,460
N (enr. NHP in 2016)	10,271	10,271	10,271	10,271
N (switch from NHP in 2017)	2,240	2,240	2,240	2,240

Notes: Table shows estimates (and standard errors in parentheses) of the proportional effect of broad-network (NHP) coverage on average insurer costs for four poisson regression specifications. Column (4) estimates the effect of NHP coverage separately by whether the patient used Partners outpatient (non-emergency) care within two years before entering the Connector, interacted with categories of distance to the nearest Partners hospital.

Figure 16: Fit of HCC Risk Score



Notes: Figure plots the fit of HCC risk score (constructed based on concurrent medical claims) with realized costs, along with line of best fit in red. Histogram bars plot the density of the HCC risk score distribution. Sample is the population of Connector enrollees 2015-2017.