Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange

Mark Shepard*

September 9, 2020

Abstract

Health insurers increasingly compete on their covered networks of medical providers. Using data from Massachusetts' pioneer insurance exchange, I find substantial adverse selection against plans covering the most prestigious and expensive "star" hospitals. I highlight a theoretically distinct selection channel: these plans attract consumers loyal to the star hospitals and who tend to use their high-price care when sick. Using a structural model, I show that selection creates a strong incentive to exclude star hospitals that challenges the effectiveness of standard policies. A key reason is the connection between selection and moral hazard in star hospital use.

JEL codes: I11; I13; I18; L13.

Keywords: Health insurance; Adverse selection; Hospital networks; Insurance exchanges.

*Harvard University, Kennedy School of Government and NBER. Email: Mark_Shepard@hks.harvard.edu. I thank the editor and three anonymous referees for extremely helpful comments. I would also especially like to thank my advisers David Cutler, Jeffrey Liebman, and Ariel Pakes for extensive comments and support in writing this paper, and Grace McCormack, Amina Abdu, and Kendra Singh for excellent research assistance. I also thank the Massachusetts Health Connector (particularly Michael Norton, Sam Osoro, Nicole Waickman, and Marissa Woltmann) for assistance in providing and interpreting the data. For helpful feedback and discussions, I thank Katherine Baicker, Amitabh Chandra, Jeffrey Clemens, Leemore Dafny, Keith Ericson, Amy Finkelstein, Jon Gruber, Ben Handel, Nathan Hendren, Kate Ho, Sonia Jaffe, Tim Layton, Robin Lee, Greg Lewis, Tom McGuire, Joe Newhouse, Daria Pelech, Jim Rebitzer, Amanda Starc, Karen Stockley, Rich Sweeney, Jacob Wallace, Tom Wollmann, Ali Yurukoglu, and seminar participants at Boston College, the Boston Fed, CBO, Columbia, Chicago Booth, Duke, Harvard, NBER (Health Care), Northwestern, UPenn, Princeton, Stanford GSB, UCLA, UCSD, and Wash U St. Louis. I gratefully acknowledge data funding from Harvard's Lab for Economic Applications and Policy, and Ph.D. funding support from National Institute on Aging Grant No. T32-AG000186 (via the National Bureau of Economic Research), the Rumsfeld Foundation, and the National Science Foundation Graduate Research Fellowship.

1 Introduction

Health insurers increasingly compete on their network of covered medical providers. Rather than cover all physicians and hospitals, insurers restrict coverage to a subset of providers with whom they have negotiated contracts. While networks have long been a factor in health insurance, "narrow" networks have recently proliferated in market-based public programs like the Affordable Care Act (ACA), Medicaid managed care, and Medicare Advantage that let individuals choose among competing plans.[1] Much more so than in employer health insurance, this structure allows for individual choice and insurer competition. But it also means that network competition may be influenced by "cream skimming" incentives associated with adverse selection (Rothschild and Stiglitz, 1976).

Although this is a classic theoretical result, whether and how selection influences insurers' incentives in setting provider networks is not well understood. While there is a large literature on adverse selection, most of it studies its impact on prices given fixed contracts, with less work on benefit competition.[2] Within the selection literature, there is no direct evidence on the connection between networks and selection incentives.[3] Most of the recent literature on narrow networks instead focuses on either measuring their cost impact (Gruber and McKnight, 2016) or on modeling their role in hospital-insurer bargaining (Ho and Lee, 2019; Liebman, 2016; Ghili, 2016).

In this paper, I study the role of selection when insurers compete on a key aspect of network quality: whether to cover the top-ranked, "star" hospitals in a market. Star hospitals (which are almost always academic medical centers) are a pervasive feature of the U.S. health care landscape. They tend to share two characteristics. First, they have reputations as centers of advanced medical treatment and research – e.g., as reflected in *U.S. News & World Report's* "Best Hospitals" rankings. Second, they tend to be quite expensive – both because they deliver more intensive services (Newhouse, 2003) and because they command high prices from insurers (Ho, 2009). As such, insurers' motives for excluding them from network may involve a mix of cost-cutting and selection. While star hospitals are often seen as "must-cover" hospitals in employer insurance, they are regularly excluded from networks in the ACA insurance exchanges (McKinsey, 2017). Understanding the reasons for this pattern is important for interpreting this trend both in the ACA and health insurance markets more generally.

To provide evidence on this issue, I draw on administrative data from Massachusetts' pioneer health insurance exchange, a model market for the ACA. Using variation in plan coverage of the state's top star hospital system, I show evidence of substantial adverse selection incentives to exclude the star providers. These incentives persist despite the use of sophisticated risk adjustment intended to neutralize adverse selection. Investigating the channels for selection, however, I argue that the finding

[1] See McKinsey (2017) and Dafny et al. (2017) on ACA exchanges; Jacobson et al. (2016) and Feyman et al. (2019) on Medicare Advantage; and Decker (2013) and Wallace et al. (2020) on Medicaid. Narrow networks are still rare in employer-sponsored insurance, but tiered networks have grown in prevalence (Claxton et al., 2019).

[2] See Einav, Finkelstein, and Levin (2010b) and Geruso and Layton (2017) for reviews of the selection literature. Some exceptions studying benefit competition include Einav, Jenkins, and Levin (2012) on credit markets; recent work on prescription drug coverage (Carey, 2017; Lavetti and Simon, 2018; Geruso et al., 2019); and work on switching rules in Medicare (Decarolis and Guglielmo, 2017). In addition, Veiga and Weyl (2016) and Azevedo and Gottlieb (2017) present theoretical frameworks for benefit determination in selection markets.

[3] The literature has focused on selection between plans with higher vs. lower cost-sharing and between HMOs and traditional (FFS) plans (see Glied (2000) and Breyer, Bundorf, and Pauly (2011) for reviews). HMOs often have narrower networks than FFS plans but also differ in a variety of other managed care restrictions.

does not immediately imply the need to modify risk adjustment or subsidize star hospital coverage. Rather, most of the selection comes through patients' incremental costs of care from the expensive star providers, creating a link between selection and moral hazard that challenges standard policy.

The paper has three main contributions. The first is the basic empirical finding of adverse selection on network generosity. I show this using data from Massachusetts' subsidized health insurance exchange (Commonwealth Care, or "CommCare"), a market established in 2006 to provide coverage for low-income individuals.[4] This setting is ideally suited to this analysis because plan financial benefits (cost sharing and covered services) are standardized, letting me study selection across plans that are nearly identical except for their provider networks. Moreover, there is variation in coverage of the state's top star hospital system: Partners Healthcare, which includes two nationally top-ranked hospitals (Mass. General Hospital (MGH) and Brigham & Women's Hospital) plus five other hospitals and more than 1,000 primary care physicians. Consistent with the star provider paradigm, Partners is a leader in medical teaching and research and is known to be quite expensive (Coakley, 2010).

The main source of evidence is a 2012 change in which a large plan drops Partners as part of shifting towards a narrow-network, low-price strategy. I use this as a natural experiment to provide reduced form evidence of selection and the mechanisms involved. Just after the exclusion, the plan sees a large exodus of high-cost consumers who live nearby a Partners hospital and/or who received care from a Partners provider during 2011. About 45% of Partners patients switch out of the plan in 2012 – a more than six-fold increase over surrounding years and a strikingly high switching rate in a setting known for inertia (Handel, 2013; Ericson, 2014). Switchers out of the plan had prior-year (2011) costs 108% higher than stayers in raw terms and 60% higher after risk adjustment – levels that made them unprofitable based on medical costs alone. Meanwhile, the plan also benefited from an influx of low-cost consumers attracted by the plan's lower price. This switching pattern illustrates the competitive logic of adverse selection. Dropping the star hospitals led many people to leave the plan, but this actually improved its bottom line – while raising costs for rivals – because the switchers were high-cost, unprofitable enrollees.

Although these findings come from a single natural experiment, the basic idea that an insurer covering (excluding) a medical provider attracts (repels) consumers with stronger demand for that provider is natural and appears generalizable.[5] Consistent with generalizability, I find similar changes in plan selection patterns for both switching and new enrollee choices, for patients of the other (non-Partners) hospitals dropped in 2012, and for Partners patients after another plan excludes the system in 2014. However, what differs across settings and providers is whether their patients have high costs. Patients of the star Partners hospitals tend to be quite expensive (even after risk adjustment), which creates the strong selection incentives to exclude Partners.

[4] This setting is distinct from Massachusetts' unsubsidized exchange ("CommChoice") studied by Ericson and Starc (2013, 2015a, 2015b). Other past work on the subsidized exchange studies the price elasticity of consumers' insurance choices (Chan and Gruber, 2010), the effects of cost sharing changes on utilization (Chandra et al., 2014), and the impact of subsidies on extensive margin consumer enrollment (Finkelstein et al., 2019). Jaffe and Shepard (2020) estimate a structural model using CommCare data to study the effects of price-linked subsidies.

[5] Indeed, this is a central assumption in the "option demand" framework of Capps, Dranove, and Satterthwaite (2003) that is the foundation of most empirical models of hospital-insurer bargaining (e.g., Gowrisankaran et al., 2015; Ho and Lee, 2017).

Why do these patients have high costs that risk adjustment fails to capture? And how should we think about the proper policy responses? My paper's second contribution is to use a simple model to conceptually analyze the sources of adverse selection related to coverage of expensive providers. I argue that selection on networks is conceptually more complex – and therefore harder to address with policy – than in standard analyses of risk selection.

My central conceptual point is that consumers may incur high spending in a plan that covers an expensive provider for two reasons, or along *two cost dimensions*. The standard dimension is *medical risk* (sickness), which leads to greater need for health care and higher costs. Most analyses of adverse selection implicitly assume that medical risk is the main or only reason for cost variation. But consumers may also vary along a second dimension: whether and how much they *choose the expensive provider for care* when sick. Choice of provider matters because treatment intensity and negotiated prices vary widely across providers (Cooper et al., 2019), and insurers cover the bulk of these cost differences rather than passing them onto patients. Moreover, people vary in their preferences for different hospitals/doctors for a host of reasons besides pure medical risk – including distance (convenience), varying perceptions of quality, and existing relationships with providers. As a result, patients with the same medical condition vary in their choice of provider and the costs they incur to the insurer.

The key economic feature of this second dimension of cost variation is that it implies *heterogeneity in consumers' moral hazard* (or costs increases) when a plan covers the expensive hospital, and thus in "selection on moral hazard" (Einav et al., 2013). While sickness is largely exogenous and results high costs in any plan, preferences for using a star hospital need not imply high costs unless the plan also covers the hospital. When the plan does so, consumers benefit from additional choice but incur higher costs if they shift to using the star hospitals' expensive care instead of cheaper providers. This is a form of moral hazard: a behavioral response to more generous benefits that raises insurer costs. Importantly, the *amount* of moral hazard varies across consumers based on their propensity to shift toward the star hospital. For instance, a consumer who lives hundreds of miles away may rarely use the star hospital and incur negligibly higher costs, while a consumer who lives nearby or has an existing relationship with a star hospital's physician may incur substantially higher costs. Insurers, however, may not charge the latter individual a higher premium, nor is risk adjustment likely to compensate for their higher costs that come from provider choices, not simply sickness.[6]

I analyze the conceptual properties of selection along these two cost dimensions in Section 2. In some ways they are similar. Both imply higher insurer average costs and (if not offset by risk adjustment) discourage coverage of an expensive hospital. This disincentive may result in either narrower networks or in pressure on star hospitals to lower prices in bargaining with insurers – with the latter being an interesting way that adverse selection can discipline star hospitals' market power.

However, in other ways selection on provider preferences/choices is different. First, it is less likely to be offset by risk adjustment, since it is not mainly about medical risk. Second, it is especially

[6] Of course, sickness is one *determinant* of provider choices/preferences, and in my empirical analysis I find that patients who choose the star hospitals are generally sicker. But the key point is that medical risk is not the *only* reason choices vary; they also vary for non-medical reasons. As a result, there is residual variation in propensity to choose the star hospitals conditional on medical risk.

likely to arise because the same preferences that lead to high costs (by using star hospitals) also drive insurance plan choice. This creates the link between insurance costs and demand that is the hallmark of adverse selection. Third, the connection between adverse selection and moral hazard (and thus, *selection on moral hazard*) makes policy responses challenging. Instead of being a simple problem to be "fixed" with subsidies or mandates, adverse selection is directly tied up in the difficult tradeoff between generous coverage and moral hazard.

My paper's third contribution is to empirically test these ideas through further reduced form analysis and a structural model. The evidence from the 2012 exclusion of Partners Healthcare is consistent with adverse selection on *both* cost dimensions. Switchers out of the plan are sicker – both on "observed" risk captured by risk adjustment and a measure of "unobserved" risk – and switching is more common among older and sicker people even controlling for existing provider relationships. However, proxies for preferring Partners (or the other dropped hospitals) based on distance and prior use of affiliated physicians are by far the strongest predictors of switching. For instance, controlling for demographics and medical risk, switching odds are *23 times* higher for Partners patients and *12 times* higher for patients of the other dropped hospitals.

Moreover, individuals who switch plans incur higher costs both because they are sicker and because conditional on risk, they have higher utilization (quantity) and use higher-price providers. A key piece of evidence comes from examining cost changes for "stayers" in the plan that narrows its network, which lets me estimate the moral hazard effect of the network change. I find a sharp 15% cost reduction for stayers at the start of 2012, with reductions occurring through both lower quantity and price of care. Cost reductions are much larger for Partners patients (about 30%) than for other enrollees (about 9%), consistent with the model's key prediction of heterogeneity in moral hazard.

To help interpret these results, I use the CommCare data and the network change to estimate a structural model of insurance plan demand and costs. The setup is standard and consists of two pieces: (1) a plan choice model, capturing how consumers trade off plan premiums vs. networks, and (2) a cost model, estimated using the claims data. Relative to past work, the main innovation is to allow for more detailed heterogeneity and use the micro data to better capture the correlations between costs and preferences for providers, which are critical for selection patterns. I pay special attention to the identification of the key parameters, using within-plan premium variation from subsidies to identify consumer price sensitivity and using the 2012 network change to identify the cost effect of the narrower network.

The estimates suggest that while enrollees significantly value better networks – and especially Partners coverage – their willingness to pay (WTP) is lower than has been found for higher-income populations (Ericson and Starc, 2015b) and falls well short of the cost savings (moral hazard reduction) from excluding the star hospitals. For instance, the average WTP for the broader network among Network Health's 2011 enrollees is just \$11 per month versus average cost savings of \$58 per month. Even more challenging is that most of the correlation between WTP and high risk-adjusted costs comes through larger moral hazard for high-WTP types (i.e., "selection on moral hazard"). The incremental (or "marginal") cost curve for the broader network is steeply downward sloping and more than three times above the WTP curve throughout the distribution.[7] Thus, any policies to subsidize or mandate star hospital coverage would need to be motivated by rationales other than standard market efficiency, including equity or externalities of their teaching/research mission.

These estimates indicate the challenge of using insurance plan choice to allocate access to star hospitals. A central problem is that a single premium cannot efficiently sort consumers who incur very different incremental costs from star hospital coverage (Bundorf, Levin, and Mahoney, 2012). Indeed, I find that the marginal cost curve is steeper than WTP, creating the conditions for "backward sorting" highlighted in recent work (Marone and Sabety, 2019). The most natural policy responses involve not simply "fixing" adverse selection but rather addressing the moral hazard directly. For instance, regulators could allow insurers to charge patients higher "tiered" copays at expensive providers (Prager, 2018), something that was disallowed in the Massachusetts exchange. But this policy would need to be carefully weighed against losses in risk protection, especially for the low-income population being served.

An important limitation of my results is that they take a "static" view of consumer demand and insurer cost structure for Partners coverage, holding them fixed as estimated in the data. I do not attempt to simulate long-run equilibrium but instead view the results as shedding light on an important economic force influencing market outcomes. For costs, adverse selection improves insurers' threat point in negotiations with star hospitals, which may lead either to network exclusion (as occurred in Massachusetts) or to star hospitals accepting lower prices. For demand, I find evidence that consumer preferences for Partners are partly (though not entirely) driven by patient-provider relationships based on recent past use, creating a form of state dependence. While I do not have enough information to model this complex dynamic relationship, state dependence suggests that long-run consumer welfare losses from narrower networks may be less than the short-run losses in my static estimates.[8]

The results in this paper are important for several reasons. First, they show the continued relevance of adverse selection, even in insurance markets that seek to address it through regulation and risk adjustment.[9] They suggest a general theoretical channel – preferences for using high-cost providers – through which selection is likely to persist. Second, they illustrate the powerful economic forces pushing towards narrower networks in individual health insurance markets like the ACA exchanges. Finally, they show that the close link between selection and moral hazard makes policy responses challenging. Selection on moral hazard is no longer a technical problem that can be "fixed" with smarter risk adjustment or subsidies (Einav et al., 2016); rather it is an economic problem that runs up against the fundamental insurance tradeoff between risk protection and moral hazard.

The paper proceeds as follows. Section 2 presents a model formalizing the paper's main ideas. Section 3 introduces the Massachusetts exchange setting and data. Sections 4-5 show reduced form

[7] Although part of these marginal costs reflect the star hospitals' high prices (which partly reflect markups), I find that WTP is still well below "adjusted" marginal cost curves that apply reductions to Partners prices of up to 50%.

[8] See Raval and Rosenbaum (2018) for evidence that *both* state dependence and heterogeneity drive hospital choices for births. Note that while consumer welfare losses may diminish over time, cost savings will also likely decline, as patients who find a new doctor no longer tend to use the star hospitals' expensive care. Therefore, the *net* cost/benefit of star hospital coverage could be either larger or smaller in the long-run.

[9] This finding contributes to an applied theory literature on "service-level selection" (Frank et al., 2000; Ellis and McGuire, 2007), which has not previously studied networks as a selection tool.

evidence of adverse selection and analyze the mechanisms involved. Section 6 presents the structural model and analyzes its estimates. The final section concludes.

2 Model and Empirical Predictions

I start with a model to formalize the mechanisms for adverse selection on provider networks. In general, adverse selection occurs when consumers with high costs sort into generous insurance plans. Given restrictions on price discriminating across consumers, this sorting pushes up premiums and gives insurers a greater incentive to cut back on benefit generosity. This standard logic is general and still applies in my model. What my model highlights is that the reasons for being a "high-cost consumer" are more complex than is typically appreciated. In addition to *medical risk* – the standard channel for cost variation – consumers vary in their *propensity to choose a high-price provider for care* when sick. This is relevant because hospital prices and treatment intensities vary substantially (Cooper et al., 2019), with star hospitals having some of the highest costs (Ho, 2009). Consumers who prefer using a star hospital both incur high (risk-adjusted) costs and demand plans that cover their preferred hospital, creating selection incentives against plans that cover the star hospitals.

In what follows, I present a model that formalizes these two cost dimensions (Section 2.1), shows how this leads to selection incentives for network design not addressed by risk adjustment (Section 2.2), and discusses the complex welfare and policy tradeoffs (Section 2.3).

2.1 Model Setup and Two Cost Dimensions

Consider an insurance market where single-plan insurers ($j \in \{1, ..., J\}$) compete on premiums and provider networks. Although I do not solve for the full model, we can think of competition occurring in a three-step process. First, networks (N_j) are negotiated between insurers and providers (h) and (if agreement occurs) characterized by bilateral hospital-insurer prices ($\{\tau_{jh}\}_{j,h \in N_j}$), whose structure I detail further below. Second, insurers observe networks and compete on plan premiums (P_j). Finally, consumers (i) choose among plans based on networks and subsidized premiums ($P_{ij}^{cons} = P_j - S_{ij}$) and while insured, experience medical shocks and incur costs.

The core of my paper is empirical evidence of adverse selection on star hospital coverage. To analyze this core finding, I isolate a key part of this larger game: a single insurer j's decision whether to cover a set of star providers (h^S) at a given set of hospital prices (τ_{jh^S}), conditional on all other networks and payment rates (both for j and other insurers).[10] Specifically, assume insurer j chooses between a narrower network, N_j^0, that excludes h^S and broader network, $N_j^1 = N_j^0 \cup h^S$. Denote plan j's choice with the binary variable $n \in \{0, 1\}$. Based on this decision, all plans can adjust premiums (in stage 2 of the game) – denoted as $P_j(n)$ and $P_{-j}(n)$ – and consumers choose plans (stage 3). The profitability of plan j's decision depends on its effects on demand and costs, which I cover in turn.

[10] While this is is a small part of the larger process, analyzing it in detail shows how selection matters for outcomes in the broader game. The setup also closely matches the natural experiment I observe in the data so is useful for generating testable predictions.

Demand Let $U_{ij}(n)$ be consumer utility (in dollars) for plan j given its coverage of the star hospital. The networks of other plans are fixed, so $U_{i,-j}(0) = U_{i,-j}(1)$. Consumers choose the plan that maximizes utility minus price: $j_i^* = \arg\max_k \{U_{ik}(n) - P_{ik}^{cons}(n)\}$. Let $\widetilde{D}_{ik}(n, P)$ be an indicator (or probability if utility is random) for whether i chooses plan k. Willingness to pay (WTP) for coverage of h^S in plan j is $\Delta WTP_i \equiv U_{ij}(1) - U_{ij}(0)$. It is natural to expect that this WTP for star hospital coverage will be larger for individuals who expect to use it for care. This is a key feature of the influential "option demand" model of Capps, Dranove, and Satterthwaite (2003) and has been borne out in empirical estimates (e.g., Ho, 2006). In general, individuals with high ΔWTP_i will be more likely to *switch* to plan j if it adds the star hospital to its network. Whether this results in adverse selection depends on how ΔWTP_i is correlated with costs.

Insurer Costs Let $C_{ij}(n)$ be the insurer's expected costs for consumer i given network $n \in \{0, 1\}$. One can think of $C_{ij}(n)$ as a "potential outcome," with j's network choice $n \in \{0, 1\}$ being the "treatment." I call $\Delta C_i \equiv C_{ij}(1) - C_{ij}(0)$ the *causal effect* on insurer costs (for consumer i) of the broader network, which I will distinguish from the *selection effect* of changing consumer sorting across plans. ΔC_i can also be thought of as a form of "moral hazard" – a change in costs due to individuals' *behavioral responses* of shifting their provider choices in response to the broader network. Because h^S is expensive, we expect that $\Delta C_i \geq 0$, though it need not always be.

A key point of my paper is that medical costs vary for two reasons, or along *two dimensions*: (1) due to varying medical risk and (2) due to varying propensity to choose expensive providers. While medical risk is exogenous, provider choices are affected by the network. To see how this plays out, suppose that individuals face risk probabilities $r_{i,d}$ of needing care for various diagnoses $d \in \{1, ..., D\}$. When ill, the individual chooses among in-network providers for care. Suppose that ω_d is the typical cost (or "severity") of illness d but that the insurer-paid cost equals $\omega_d \cdot \tau_{jh}$ if the individual chooses provider h, where τ_{jh} is the hospital-insurer price.[11] The expected cost for individual i under the narrower network is:

$$C_{ij}(0) = \sum_d \underbrace{r_{id} \cdot \omega_d}_{\text{Medical risk}} \times \underbrace{\left[\sum_h \tau_{jh} \cdot s_{idh}(0)\right]}_{\text{Price of chosen providers}} \tag{1}$$

where $s_{idh}(0)$ is the probability i chooses provider h under the narrower network.

Equation (1) formalizes how costs break down into these two dimensions. If the insurer expands to the broader network, costs increase due to moral hazard as patients shift towards the expensive star hospital. In the (simplifying) case where $s_{idh^S}(0) = 0$, moral hazard for i equals:

$$\Delta C_i = \sum_d \underbrace{r_{id} \cdot \omega_d}_{\text{Medical risk}} \times \underbrace{s_{idh^S}(1)}_{\text{Propensity to choose} h^S} \times \underbrace{\left(\tau_{j,h^S} - \overline{\tau}_{ij}^{Other}\right)}_{\text{Higher price of} h^S} \tag{2}$$

[11]In this setup, "prices" τ_{jh} capture high insurer costs *either* because of an intensive treatment style (high quantity of care) or because of high negotiated payment rates per treatment. For many purposes, these two factors enter identically. I discuss the distinction further below and separate the two in my empirical work using a price-quantity decomposition (Section 5.1). To be clear, the model allows high-price providers to be higher quality, as reflected in consumer utility. But the key fact is that *insurers* (not patients) pay for this higher fee, which will sometimes lead to overuse.

where $\overline{\tau}_{ij}$ is the average costliness of the providers i substitutes away from to use the star hospital.[12] Notice that moral hazard is unlikely to be uniform across people. It will instead be higher for people with higher risk and greater propensity to choose the star hospital when covered. These are precisely the groups we expect to have high WTP for a plan covering the star hospitals.

Risk Adjustment and Insurer Profits Although consumer costs vary, the exchange attempts to offset this with risk adjustment. Let $\varphi_i = E\left(C_{ij}|Z_i\right)/\overline{C}$ be consumer i's risk score, designed to predict their relative costliness based on observables Z_i. While insurers must charge the same price P_j to everyone, the regulator adjusts payments so the insurer receives revenue $\varphi_i P_j$ for consumer i. The consumer's profitability therefore equals $\varphi_i P_j(n) - C_{ij}(n)$ under network $n \in \{0,1\}$. Following Curto et al. (2014), it is useful to factor out φ_i and write profits as:

$$\pi_j(n,P) = \sum_i \left[P_j - \underbrace{C_{ij}^{RA}(n)}_{\text{Risk-adj costs}} \right] \cdot \underbrace{D_{ij}(n,P)}_{\text{Scaled demand}} \tag{3}$$

where $D_{ij}(n,P) = \varphi_i \widetilde{D}_{ij}(n,P)$ is risk-scaled demand, and $C_{ij}^{RA}(n) \equiv C_{ij}(n)/\varphi_i$ is risk-adjusted costs. Profits can also be written in terms of total demand and risk-adjusted average costs:

$$\pi_j(n,P) = \left[P_j - AC_j^{RA}(n,P)\right] \cdot D_j(n,P) \tag{4}$$

where $D_j(n,P)$ is total scaled demand and $AC_j^{RA}(n,P) \equiv \frac{1}{D_j}\sum_i C_{ij}^{RA}(n) D_{ij}(n,P)$ is risk-adjusted average costs. Note that an insurer breaks even only if $P_j \geq AC_j^{RA}(n,P)$.

2.2 Selection Incentive and the Limits of Risk Adjustment

Now consider the insurer's profitability of shifting from the narrower to broader network, $\Delta\pi_j \equiv \pi_j\left(1,P^1\right) - \pi_j\left(0,P^0\right)$, where superscript variables refer to outcomes under network scenario $n \in \{0,1\}$ and "Δ" variables refer to the change between the two scenarios. The change in profits can be decomposed as:

$$\Delta\pi_j = \underbrace{\left(P^0 - AC_j^{RA,0}\right) \cdot \Delta D_j}_{\text{Demand change (fixed markup)}} + \underbrace{\left(\Delta P_j - \Delta AC_j^{RA}\right)}_{\text{Change in markup}} \cdot D_j^1 \tag{5}$$

The change in profits depends on the change in demand at a pre-existing markup and the effect on the insurer's markup. The latter in turn depends on how much the insurer can raise prices and how the change affects its (risk-adjusted) average costs. The change in average costs, in turn, can be decomposed as:

$$\Delta AC^{RA} = \underbrace{\overline{\Delta C}}_{\text{Moral hazard (avg)}} + \frac{1}{D_j^1}\sum_i \left[\underbrace{\left(C_{ij}^{RA}(0) - AC_j^{RA,0}\right) \cdot \Delta D_{ij}}_{\text{(1) Selection on baseline costs}} + \underbrace{\left(\Delta C_i^{RA} - \overline{\Delta C}\right) \cdot D_{ij}^1}_{\text{(2) Selection on moral hazard}} \right] \tag{6}$$

[12] Mathematically, $\overline{\tau}_{ij}^{Other} \equiv \left(\sum_d s_{idh^S}(1)\right)^{-1} \sum_d \sum_{h \neq h^S} \tau_{jh} \cdot \left(s_{idh}(1) - s_{idh}(0)\right)$.

The change in average costs equals the moral hazard effect of the broader network for an average consumer ($\overline{\Delta C}$) plus two "selection" terms: (1) selection on baseline costs, and (2) selection on moral hazard. Term (1) is the traditional way of viewing adverse selection (e.g., in models without moral hazard). There is an *adverse selection incentive* if when the plan broadens its network, people who join the plan ($\Delta D_{ij} > 0$) have above-average risk-adjusted costs and/or if those who leave the plan ($\Delta D_{ij} < 0$) have low risk-adjusted costs. Term (2) captures *selection on moral hazard*: whether individuals who choose the broad-network plan have especially large incremental costs due to the expanded network. This is a hybrid term that involves both selection (the plan attracting high-cost individuals, which we usually want to offset with risk adjustment) and moral hazard (real increases in resource use, which we typically do *not* want to offset).

Equations (5) and (6) yield two insights. First, they provide a way to empirically test for adverse selection, as I discuss below. Second, they make clear that *if* there is an adverse selection incentive, it will reduce the profitability of a broader network. But this raises a question: why would there be a large selection incentive, especially with risk adjustment in place to try to offset it?

The insights from the the previous subsection suggest a natural answer. Consumer costs vary for two reasons: medical risk and propensity to choose higher-price providers. But risk scores, φ_i, are based only on medical risk factors like age, sex, and diagnoses. Even if risk adjustment were "perfect" – that is perfectly able to predict illness risks r_{id} – it would be missing an entire dimension of costs. Of course, risk adjustment may also be imperfect at predicting medical risk, with unobserved risk also correlated with WTP for the broader network. Both dimensions of cost variation can contribute to the selection incentive.

These observations raise a question: should we simply add predictors of provider choices (e.g., location) to the observables used to calculate risk scores? Unfortunately, this will not solve the problem as long as risk scores are independent of the network. Recall that risk-adjusted costs equal $C_{ij}^{RA}(n) = C_{ij}(n)/\varphi_i$. As shown above, costs change with the network in a way that varies across consumers. As long as φ_i is independent of the network, risk-adjusted costs will vary across consumers, and selection incentives will remain.[13] But if regulators set φ_i in a way that compensates insurers for certain consumers' higher costs when star hospitals are covered, they create an implicit subsidy for the broader network. While such a subsidy could be desirable, it is a policy choice with clear tradeoffs, not a mere technical fix to risk adjustment.

Underlying these complications is the reality that adverse selection in this setting is closely linked to heterogeneous moral hazard driven by high star hospital prices. In equation (6), the correlation between demand for star hospital coverage (ΔD_{ij}) and costs is driven both by "cost levels" ($C_{ij}^{RA}(0)$) and by "cost changes" (ΔC_i^{RA}). This creates a situation of "selection on moral hazard" (Einav et al., 2013), which has been shown to create a particular challenge for risk adjustment and to lead to complex policy tradeoffs (Einav et al., 2016). In a mathematical sense, a single risk score φ_i cannot accurately capture variation in both $C_{ij}(0)$ and $C_{ij}(1)$ when ΔC_i varies separately from $C_{ij}(0)$. In an

[13] With multiplicative risk adjustment as in this model (as opposed to additive risk adjustment in some other settings like the ACA), this statement technically requires that *proportional* incremental costs, $\Delta C_{ij}/C_{ij}(0)$, vary across consumers. Inspection of equations (1) and (2) shows that this will be true as long as τ_{jh^S} is high and $s_{idh^S}(1)$ varies across consumers.

economic sense, the presence of moral hazard means that fully offsetting adverse selection may not be optimal because it encourages overuse of the star hospitals. There can even be cases of "backward sorting" where people with the highest demand for the generous network should not get it because their incremental costs are larger than their willingness to pay (Bundorf et al., 2012; Marone and Sabety, 2019).

Empirical Tests for Selection Equation (6) motivates my empirical strategy that examines *changes in plan choices* (e.g., from plan switching behavior) following a network change and how these correlate with individual-level costs prior to the change. Empirically, I look at a network narrowing – in the above notation, a shift from N_j^1 to N_j^0 – so the signs of demand changes are reversed but the idea is the same. If people who leave the plan in response to a narrower network ($\Delta D_{ij} > 0$ in the notation above) are high-cost and/or those who join ($\Delta D_{ij} < 0$) are low-cost, the insurer will have a selection incentive against the broader network. This prediction is testable given data on choices and costs before and after the network change. Notice that this test can be thought of as first-differences version of the classic positive correlation test (Chiappori and Salanie, 2000): it asks whether a plan that *changes its network* in turn attracts a *changing selection* of consumers.

The theory above suggests two further tests to analyze the *source of high costs*. First, equation (1) provides a framework for decomposing costs, which I implement empirically in Section 5. Second, equation (2) suggests analyzing heterogeneity in incremental costs due to a narrower network, to understand the role of selection on cost levels vs. moral hazard. To implement this empirically, I examine cost changes for enrollees who stay with the plan that narrows its network and correlate these changes with measures of propensity to use the star hospitals.

2.3 Implications and the Role of State Dependence

What are the implications of adverse selection on networks through the mechanisms laid out above? Although my primary focus in this paper is explaining (and providing evidence of) these mechanisms for selection, it is worth discussing the implications for market outcomes, welfare, and policy. In doing so, I think about the broader model mentioned at the start of the section in which networks and bargained hospital prices can adjust.

Market Implications The implications for market outcomes are in some ways standard. As in most models, adverse selection pushes up the average cost and price of generous plans, leading lower-demand enrollees to select out of them. Strong enough adverse selection may lead to "unraveling" or elimination of generous plans from the market (e.g., Rothschild and Stiglitz, 1976; Cutler and Reber, 1998; Azevedo and Gottlieb, 2017).

Less standard is the implications for hospital-insurer bargaining. Adverse selection may *discipline hospital market power*. To see this, note that the profitability of excluding a star hospital ($\Delta\pi_j$ in equation (5)) is a key term in hospital-insurer Nash bargaining models. Adverse selection improves the insurer's threat point in bargaining with a hospital, since the insurer knows it will benefit from selection if it drops the hospital. This force is relevant for bargaining models in settings that allow for consumer plan choice (e.g., Ho and Lee, 2017). It represents a distinction from employer insurance settings that typically do not include choice among multiple insurers (e.g., Gowrisankaran et al., 2015) and may be one reason narrow networks are less common in job-based insurance (Claxton et al., 2019). Importantly, because hospital prices can adjust down – which diminishes (or even reverses) the hospital price/choice dimension of costs – adverse selection on star hospital coverage need not lead to unraveling but may simply reduce hospital prices. What actually occurs will depend on the star hospitals' cost structure and willingness to accept lower payments.

Welfare and Policy Implications The welfare implications of adverse selection on networks of course depends on the market outcomes. But suppose that insurers *do* drop the star hospitals and/or plans covering them become expensive. Would this represent a loss for consumer and social welfare? And what are the policy implications? There are several reasons the answer is less obvious than in a textbook model of adverse selection.

First, as discussed above, the link between selection and moral hazard means that simple fixes to selection incentives like a mandate may not be optimal. Patients may overuse expensive hospitals when insurance covers their higher costs, so extending universal access to star hospitals may be wasteful. On the other hand, allowing the market to unravel may deprive the sickest patients who need the star hospital's specialized care. Mandates and risk adjustment are not well-suited to addressing the challenges arising from heterogeneous use of the star hospitals. Instead, policies that directly target the utilization tradeoff – like tiered copays for expensive providers (Prager, 2018) or physician incentives to consider costs when making referrals (Song et al., 2012; Ho and Pakes, 2014) – may be fruitful.

Second (and in the opposing direction), providing access to star hospitals even at high prices may be *socially valuable* for reasons that go beyond standard market welfare metrics. High academic hospital prices may cross-subsidize teaching and research, which likely have positive externalities. Ensuring access to the best hospitals among low-income enrollees in ACA and Medicaid markets may also be important for social equity. Ultimately, a key take-away of the paper is that because of adverse selection, "letting the market work" will tend to squeeze star academic hospitals via narrower networks and less bargaining power. Policymakers will need to make a judgment about whether they find this outcome desirable or wish to counteract it.

Finally, demand for star hospitals (which drives plan choices) may be high because of *state dependent relationships* with the hospital or its affiliated physicians, which I discuss next.

Role of State Dependent Relationships Some consumers may have high demand for star hospitals simply because they have developed a relationship with the hospital or its affiliated physicians. This means that preferences/demand for star hospitals are *state dependent*: they depend not only on static factors but on past history. I find suggestive evidence of this in my empirical work, as has other recent work (Raval and Rosenbaum, 2018; Higuera et al., 2018). Even if a patient only slightly preferred the star provider initially, over time a relationship builds and creates a cost of switching to a new provider. These switching costs may take many forms – including search costs, medical care disruption, and loss of relationship-specific knowledge – which complicates their welfare interpretation.[14]

State dependence adds a dynamic twist to the positive economics of the adverse selection channel I highlight but does not change the basic idea. Because loyalty to an expensive hospital is not simply exogenous but develops with use, an insurer that covers the star hospitals *develops* more patients who are loyal to it and who therefore incur high costs. The selection disincentive may be smaller initially and grow over time or based on outside access to the star hospitals. The basic selection idea is similar but simply varies dynamically and based on market conditions.

But state dependence can have more complex implications for welfare. Traditionally, it creates a wedge between short- and long-run consumer welfare. In the short run, a person may have very high value for a plan that covers their current hospital/doctor. But if forced to switch, the long-run welfare loss after incurring the switching cost would be less. Thus, in this sense, state dependence may tend to mean that lost access to star hospitals has less of a patient welfare cost in the long-run than in the short-run. If a two-tier medical system develops in which low-income patients never get access to star hospitals, it may be less disruptive to patient welfare than the evidence in this paper suggests. However, as the theory shows, preferences also drive *costs* through patient's provider choices. Thus, the long-run tradeoff between consumer WTP and cost for a broader network could be either more or less favorable in the long run.

3 Massachusetts Exchange Setting and Data

3.1 Setting: Massachusetts Subsidized Exchange (CommCare)

I study Massachusetts' subsidized health insurance exchange – called Commonwealth Care, or CommCare. Created in the state's 2006 "Romneycare" health reform, CommCare operated from 2006-2013 to provide subsidized coverage to low-income people (below 300% of poverty) not eligible for employer insurance or other public programs.[15] Enrollees could choose among competing private plans in a centralized marketplace. Over the 2010-2013 period I focus on, the exchange featured five competing insurers and averaged 170,000 enrollees – making it a substantial market but still only a small portion of the state's population of 6.6 million.

CommCare is a good setting to study the selection implications of provider networks (and star hospital coverage in particular) for several reasons. First, the exchange standardized essentially all benefits other than networks. By rule, each insurer offered a single plan with state-specified covered services and patient cost sharing rules.[16] This structure – which is more standardized than ACA exchanges but similar to Medicaid managed care programs – lets me study plans that differ in network but are nearly identical on other dimensions.

[14]Recent evidence on exogenous breaks to patient-doctor relationships suggests they may involve first-order medical consequences like increased ER visits and hospitalizations and reduced use of preventive care (Sabety, 2020).

[15]A separate market called "CommChoice" offered unsubsidized plans for all others (for research on CommChoice, see Ericson and Starc 2015b; 2015a; 2016). In the ACA, unsubsidized and subsidized enrollees are pooled into a single exchange, while people below 138% of poverty are eligible for Medicaid in states that have chosen to expand the program.

[16]The only exceptions to this identical coverage were: (1) prescription drug formularies for above-poverty enrollees, subject to minimum standards, and (2) a few "extra benefits" like gym memberships.

Second, like the ACA, CommCare used sophisticated policies to address risk selection. In addition to subsidies and a mandate to encourage broad participation in the market, it risk adjusted payments to insurers based on enrollee observables.[17] Specifically, the exchange used demographics and past diagnoses to assign each enrollee a "risk score" predicting their relative costliness. Risk scores multiplied the insurer's price (P_j), so the insurer received $\varphi_i P_j$ for an enrollee with risk score φ_i. While there is debate on how well risk adjustment has worked elsewhere (see Brown et al., 2014; Newhouse et al., 2015), the methods used by CommCare were state-of-the-art. The one notable limitation was the choice to use "prospective" or "*ex-ante*" risk scores based on only the prior year's claims data, whereas the ACA uses a "concurrent" risk score (the HHS-HCC method) based on current year's claims data. While *ex-ante* risk adjustment limits problems with indirectly tying risk scores to current utilization (Geruso and McGuire, 2016), it also misses information, especially for new enrollees who lack past claims data. I use the concurrent HCC score as a way of capturing medical risk unobserved by CommCare's *ex-ante* risk score.

Third, Massachusetts has a clear pair of star hospitals: Massachusetts General Hospital (MGH) and Brigham & Women's Hospital (BWH), which are owned by the Partners Healthcare System and affiliated with Harvard Medical School. *U.S. News & World Report* perennially ranks these as the top two hospitals statewide and among the top 10 nationwide. This position has given them the perception of "must-cover" hospitals that can command high prices, as has been repeatedly documented for commercial insurance (e.g., Coakley, 2010; CHIA, 2014). Further, Partners Healthcare is the state's largest health system, giving it substantial market power even beyond its star hospitals. As of 2012, it also owned five community hospitals around Boston and employed about 1,100 primary care physicians. Thus, Partners represents a pure (if perhaps extreme) example of two attributes that are known to drive high hospital prices: star status (Ho, 2009) and high market share (Cooper et al., 2019).

Table 1 confirms this pattern of high star hospital prices for CommCare, drawing on estimates from the price-quantity model in Section 5.1. The table reports price estimates for the 10 highest-price hospitals in the data. Column (1) shows raw average payments per admission (straight from the claims data), and columns (2)-(4) report model estimates of relative price and patient severity (vs. an average value of 1.0 for each). The two star Partners hospitals (MGH and Brigham & Women's) are the most expensive by a large margin. Their relative prices of about 1.60 are more than 20% above the next highest-price hospital and 60% above the average hospital.[18] The star hospitals also attract some of the sickest patients, with average severities 25-37% above average. Thus, the table illustrates the phenomenon of high prices and sicker patients for academic medical centers (AMCs) – of which the star hospitals are just the most extreme example. All six of the state's top AMCs (as designated by the state) appear in this top-10 list, and all six have above-average severity.

[17] CommCare also had a reinsurance program, which covered 75% of any enrollee's costs exceeding $150,000 per year. This very high cutoff meant reinsurance played a minor role, covering just 0.03% of enrollees and 1% of costs.

[18] A natural question is whether the star hospitals' high prices reflect higher underlying costs or higher markups. The answer appears to be both. Based on a state report of average cost per severity-adjusted admission (CHIA, 2014b), the Brigham and MGH have the highest casemix-adjusted costs of any large general acute hospital, with costs about 30-50% above average. While these costs are not perfectly comparable to the CommCare prices (since the casemix adjustment may differ), note that prices exceed the average by a similar or larger percent (~60%) than costs (30-50%). This suggests that as long as price exceeds costs, markups in dollar terms are also high at the star hospitals.

Table 1: Hospital Prices: Most Expensive Hospitals for CommCare Insurers

	Hospital	**System**	**Teaching Status**	**Raw Data** Avg. Insurer Payment (1)	**Hospital Price Model** Relative Price Estimate (2)	*Std. Err.* (3)	Rel. Patient Severity (4)
1	Brigham & Women's	Partners	AMC	$23,525	1.62	*(0.04)*	1.37
2	Mass. General (MGH)	Partners	AMC	$21,090	1.58	*(0.04)*	1.25
3	Boston Med. Ctr.	BMC	AMC	$16,478	1.29	*(0.03)*	1.20
4	Baystate Med. Ctr.	Baystate	Teaching	$13,411	1.27	*(0.03)*	0.99
5	UMass Med. Ctr.	UMass	AMC	$14,540	1.19	*(0.03)*	1.16
6	St. Vincent's	Vanguard	Teaching	$11,824	1.10	*(0.03)*	0.99
7	Southcoast Hospitals	Southcoast	---	$12,402	1.10	*(0.03)*	1.06
8	Beth Israel Deaconess	CareGroup	AMC	$12,266	1.06	*(0.03)*	1.11
9	Tufts Med. Ctr.	Tufts	AMC	$15,378	1.02	*(0.03)*	1.50
10	Carney Hospital	Steward	Teaching	$9,200	1.02	*(0.03)*	0.85
	Average Hospital	---	---	*$11,062*	*1.01*	---	*1.00*
	Non-Top 10 Hospitals	---	---	*$7,972*	*0.84*	---	*0.88*

NOTE: The table shows the 10 highest-price acute care hospitals in the CommCare data, ranked by the inpatient hospital price measure in column (2). Hospital system is as of 2013, and teaching status of "AMC" refers to academic medical centers, the six most sophisticated academic hospitals as designated by the state. All estimates are from the hospitalization dataset described in Section 3.2, limited to in-network admissions. Column (1) shows the average insurer payment per admission directly from the raw data. Columns (2)-(3) show the in-network relative price estimates (for $t = 2011$) and standard errors from the model. Column (4) reports the average severity for each hospital's patients. Both prices and severities are relative measures (mean 1.0) and are calculated from the estimates of (13) as described in the text. (Price has mean 1.01 for the average hospital in the penultimate row because the sample is restricted to in-network admissions.) Price standard errors come from calculating confidence intervals for the appropriate weighted sum of α and β coefficients, then exponentiating the upper and lower confidence bounds and dividing the resulting interval width by 2*1.96.

Finally, CommCare features substantial variation in enrollee premiums that is useful for estimating a model of insurance choices. This variation comes from two sources. First, insurers vary prices over time as they acclimate to the new market and adjust strategy. Of course, these price changes may be endogenous to plan quality shifts. Therefore, I also use a second source of variation: *subsidies* that differ by income group and that affect *premium differences* across plans. Notably, enrollees earning below 100% of poverty are fully subsidized: they pay $0 for any available plan. Higher-income enrollees get the same plans but pay more on the margin for higher-price plans. This sets up a natural identification strategy for premium coefficients in my plan demand model. Effectively, below-poverty enrollees serve as a "control group" for inferring shifts in unobserved quality, with the remaining changes in higher-income groups' demand coming from premiums. I discuss this strategy further in Section 6.1 and present statistics on the underlying premium variation in Appendix B.

3.2 Administrative Data: Plan Enrollment and Insurer Claims

I use administrative data on plan enrollment and insurance claims for all CommCare plans and enrollees from fiscal 2007-2014.[19] For each (de-identified) enrollee, I observe demographics, plan enrollment history, and claims for health care services while enrolled in the market. The claims include information on patient diagnoses, services provided, the identity of the provider, and the actual amounts the insurer paid for the care. I use the raw data to construct three analysis datasets, which I describe in turn.

Hospitalization Dataset The first dataset is for hospital choices and costs. From the claims, I pull out all inpatient admissions at general acute care hospitals in Massachusetts during fiscal years 2008-2013, the period over which I observe networks. Constructing an admission-level dataset from underlying insurance claims – which typically feature multiple claims per admission – is an involved process; I discuss the details in Appendix A.1. For each admission, I use the information on the claims to observe the treating hospital, the principal diagnosis and diagnosis-related group (MS-DRG used by Medicare), comorbidities, and total insurer payments (including both facility fees and physician professional services). To this, I add hospital characteristics from the American Hospital Association (AHA) Annual Survey and define travel distance using the driving distance from the patient's home zip code centroid to each hospital.[20] I use this dataset to estimate the hospital price model (see Section 5.1) and the hospital choice model (see Section 6).

Plan Choice and Cost Dataset The second dataset is for insurance plan choices and total costs. I construct a dataset of available plans, plan characteristics (including premium and network), and chosen options during fiscal 2008-2013. This dataset is constructed at the level of instances of enrollees making a plan choice, which occur at two times: (1) when an individual newly enrolls in CommCare (or re-enrolls after a gap), and (2) at annual open enrollment when current enrollees can switch plans. A key difference between these two situations is their default choice. New and re-enrollees must make an active choice to receive any coverage,[21] while passive current enrollees are defaulted to their current plan. For each enrollee x choice instance, I calculate insurer costs over the subsequent year (from the claims data) and add on enrollee attributes, including demographics and risk scores. I also use the claims data to decompose this cost into prices vs. quantities, as discussed in Section 5.1.

Outpatient Care Provider Use Variables I construct measures of whether enrollees have used certain hospitals or their affiliated physicians and community health centers (CHC) for outpatient care (see Appendix A.2 for details). These present a broader picture of provider utilization to understand whether a patient's access will be curtailed by the network limits. Starting from the full claims data, I exclude inpatient and emergency department care, following a similar definition as for the hospitalization dataset. I then limit to outpatient and professional services using a flag given by the

[19]The data was obtained via a data use agreement with the Massachusetts Health Connector, the exchange regulator. To protect enrollees' privacy, the data was purged of all identifying variables.

[20]I thank Amanda Starc and Keith Ericson for sharing this travel distance data.

[21]This rule had one exception. Prior to fiscal 2010, the exchange auto-assigned plans to the poorest new enrollees who failed to make an active choice. I exclude these passive enrollees from the plan choice estimation dataset.

data provider. Finally, I code the hospital or CHC (if any) at which the outpatient care was delivered using the name of the billing provider on the claims. The key variables for my analysis are whether enrollees received non-ED outpatient care via a doctor treating at a Partners hospital/CHC or another hospital excluded in the 2012 network change (which I discuss next).

Summary Statistics Appendix Table 5 shows summary statistics for these datasets. The data include 624,443 unique enrollees making 1,684,203 plan choices and experiencing 70,094 hospital admissions. The average age is 39.9, and 47% of enrollees earn less than poverty so are fully subsidized. There is substantial flow of enrollees into and out of the market – about 11,000 people per month (or 6.5% of the market) in steady state – giving me a significant population of active choosers to assist in plan demand estimation.

3.3 Star Hospital Coverage and 2012 Network Change

CommCare plans' hospital networks vary significantly, including in coverage of the star hospitals. Overall statistics on the size of hospital networks are shown in Appendix Figure 11. Here, I focus on the coverage of the star Partners hospitals, with additional details reported in Appendix Table 7. Up to 2011, three of the four Boston-area insurers covered the star Partners hospitals – Network Health, Neighborhood Health Plan (NHP), and CeltiCare (which newly entered the market in 2010). One plan – BMC HealthNet, which is vertically integrated with Boston Medical Center, a competitor hospital – did not cover Partners, and a final plan (Fallon) operated mainly in central Massachusetts and did not have a full Boston network.

My empirical work exploits a major change in Partners coverage in fiscal 2012. In 2012, the exchange introduced new rules encouraging insurers to compete more aggressively on premiums.[22] In response, Network Health and CeltiCare cut their prices sharply. Although CeltiCare already had a narrow network and low-cost structure (despite its covering Partners), Network Health needed to reduce costs to make this price cut feasible. To do so, Network Health dropped the Partners hospitals and associated physicians, plus several other less prestigious hospitals.[23] Figure 1 shows that two major shifts for Network Health followed. Panel A shows that its average enrollee cost fell sharply by 26%, from \$400 per month at the end of 2011 to \$296 at start of 2012. The exchange's risk adjustment partly offset this fall, but risk-adjusted costs also fell by 21%. Panel B shows that the share of Network Health's hospital admissions going to a Partners hospital fell by two-thirds, while Partners use rose in other plans.

These sharp changes reflects a combination of selection and causal cost reductions from the narrower

[22]There were two main policy changes. First, the exchange lowered the insurer price floor (a rule intended to ensure actuarial soundness of the insurer), which had in previous years been binding on CeltiCare and Network Health. Second, the exchange introduced new choice limits for enrollees below 100% of poverty, for whom all plans were fully subsidized (\$0 premiums). Starting in 2012, new enrollees in this group were limited to choosing one of the two cheapest plans, which encouraged insurers to cut prices to be one of these limited choice options.

[23]These other hospitals included one less prestigious academic medical center (Tufts Hospital), one teaching hospital (St. Vincent's in Worcester), and six community hospitals. The plan did retain two small and isolated Partners hospitals on the islands of Nantucket and Martha's Vineyard but dropped all other Partners providers.

Figure 1: Changes for Network Health around 2012 Network Change

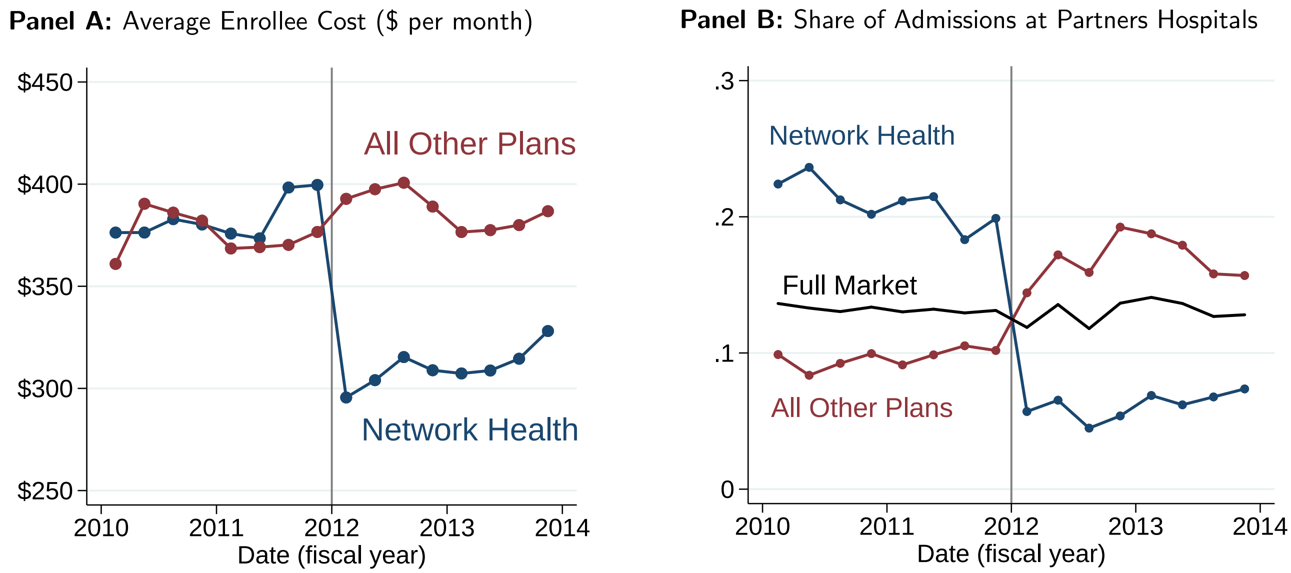

NOTE: The figures show average enrollee cost per month (left graph) and Partners hospital use shares (right graph) by enrollees in Network Health and all other CommCare plans. Each point is a quarterly average, and the vertical line marks the point where Network Health drops Partners from its network. Importantly, these patterns represent the combined effect of selection (enrollees shifting between plans) and causal effects of the change. Average costs fall sharply for Network Health at the start of 2012 (by about 25-30%), while rising somewhat in other plans. The share of admissions at Partners hospitals falls by about two-thirds for Network Health in 2012, while rising sharply in all other plans. The rise in Partners use in other plans (whose networks did not change) is consistent with the paper's main selection story: enrollees who want to use Partners shift from Network Health to other plans that cover it to facilitate this hospital choice.

network. A key goal of my analysis will be to separate out the two. One indication that selection matters is that other plans' average costs and Partners admissions *rose* in 2012, despite no major change in their networks. In particular, the two other plans covering Partners (CeltiCare and NHP) received over 90% of consumers who left Network Health in 2012. These two plans' costs and Partners use rates rose particularly sharply starting in 2012. Interestingly, Partners' market-wide share of inpatient admissions (black line in panel B) was flat through 2012, suggesting that the enrollees who most wanted Partners were able to retain access by switching plans.

After seeing higher costs in 2012-2013, CeltiCare dropped Partners in fiscal 2014, explicitly citing adverse selection as a rationale.[24] My ability to study this change is more limited because it occurs at the tail end of my data (e.g., claims data for 2014 are incomplete), but I use it for some robustness checks on the main selection findings.

By the start of the ACA in January 2014, the only plan still covering Partners was NHP, which Partners had acquired as a subsidiary during fiscal 2013. NHP's status as the only plan to cover Partners has continued through at least 2019 in the state's "ConnectorCare" program – the post-ACA

[24]In testimony to the Mass. Health Policy Commission (HPC 2013), CeltiCare's CEO wrote: "For the contract year 2012, Network Health Plan removed Partners hospital system and their PCPs from their covered network. As a result, the CeltiCare membership with a Partners PCP increased 57.9%. CeltiCare's members with a Partner's PCP were a higher acuity population and sought treatment at high cost facilities. . . . A mutual decision was made to terminate the relationship with BWH [Brigham & Women's] and MGH PCPs as of July 1, 2013."

successor of CommCare that covers the large majority of exchange enrollees. Moreover, NHP has experienced significant financial challenges, losing over \$100 million during 2014 and subsequently having to raise its price substantially. By 2019, NHP's exchange market share had fallen into single digits. Similar patterns of near-unraveling of Partners coverage have also extended to the state's Medicaid program, which contracts with most of the same insurers.[25]

4 Reduced Form Evidence of Adverse Selection

This section provides reduced form evidence of adverse selection on star hospital coverage consistent with the mechanisms shown in the theory in Section 2. To do so, I draw on the natural experiment created by Network Health dropping the star Partners hospitals (and several less prestigious hospitals) in 2012, as described in Section 3.3. I use the natural experiment to test the model's prediction that dropping the star hospitals should result in favorable selection (high-cost individuals leaving the plan) driven by individuals with high demand for the star hospitals.

To test for selection, I examine how *changes in consumer plan choices* following the network narrowing correlate with consumer costs. This can be thought of as first-differences version of the classic positive correlation test (Chiappori and Salanie, 2000): it asks whether a plan that *changes its network* in turn attracts a *changing selection* of consumers. The assumption throughout is that changing plan choices are caused by Network Health's narrower network and lower premium, and not other contemporaneous shocks – an assumption that seems reasonable given the market environment and that I provide evidence to support. Plan choice changes come in two ways: (1) through plan switching by current enrollees and (2) through different initial plan choices by new enrollees. My main analysis focuses on plan switching, which lets me study within-person demand changes and also to measure costs and other variables prior to the network change (to avoid conflating selection with causal effects of the narrower network). A downside of studying switching, however, is that it is subject to substantial inertia, which may dampen demand changes. Therefore, along with other robustness checks, Section 4.2 studies how initial plan choices change among repeated cross-sections of new enrollees. Finally, Section 4.3 considers *why* certain consumers – especially existing star hospital patients – have strong demand for these providers, especially the role of state dependence vs. persistent preference heterogeneity.

4.1 Evidence from Plan Switching

I start by examining plan switching behavior at annual open enrollment. For this analysis, I limit the sample to continuous enrollees in CommCare from the final month of year $t-1$ (when consumers have the opportunity to switch) to the start of year t (when switches take effect). I define the "switching out rate" for a plan-year (e.g., Network Health in 2012) as the number of people who switched out divided by the total who could have switched out. The "switching in rate" is defined as the number of

[25]Network Health dropped Partners in Medicaid as of the start of 2014, leaving NHP as the only managed care plan covering Partners. NHP subsequently faced large financial losses and suspended new Medicaid enrollment as of late 2016. This lasted until a major Medicaid reorganization under an Accountable Care Organization program in 2018.

switchers *into* the plan divided by the same denominator, which allows for comparing the two figures in levels.

Panel A of Figure 2 shows switching rates for Network Health in each year from 2009-2014. At the start of 2012 when its narrower network (and lower price) took effect, the plan experienced a spike in switching – to 11.2% for switching out and 7.6% for switching in. While low in absolute terms (consistent with the presence of inertia), these rates are more than double those of adjacent years.[26] This is consistent with the shift to a narrower network and lower price spurring significant changes in plan choices (i.e., ΔD_i), which is necessary for selection incentives to be relevant.

Panel B of Figure 2 shows that 2012 switches were correlated with prior-year costs in a way consistent with adverse selection. Switchers out in 2012 represent a clear outlier relative to other years when they have similar or lower costs than stayers. By contrast, 2012 switchers out have costs 108% higher than stayers ($675 vs. $324 per month). CommCare's risk adjustment (shown in open points[27]) narrows this cost gap to 60% ($508 vs. $317 per month) but does not close it. Indeed, the risk-adjusted costs of switchers out greatly exceeded the plan's price ($423 in 2011 and $360 in 2012), indicating that they were unprofitable based on medical costs alone. By contrast, switchers in for 2012 were relatively low-cost, with raw (risk-adjusted) costs 29% (20%) below stayers.

The results thus far indicate overall switching patterns consistent with adverse selection. Figure 3 shows evidence of the connection between switching out and demand for Partners and the other excluded hospitals, a key model prediction. Panel A shows that the 2012 spike was concentrated among people living closer to a Partners hospital, consistent with distance as a driver of provider choice. Aside from 2012, switching rates are low (~5-10%) and differ little based on location. But in 2012, switching spikes to 22% for people within 5 miles of a Partners hospital, versus a relatively steady 5% rate for those >25 miles away. Similarly, Panel B shows that switching was concentrated among people who had used a dropped hospital for outpatient care over the past year, a revealed preference indicator of demand. For prior-year Partners patients, the switching out rate spikes to 45% – a more than *six-fold increase* over surrounding years. Switching also jumps to 24% for patients of other dropped hospitals. By contrast, switching for all other enrollees was much lower (3%) and essentially flat versus adjacent years.[28]

These patterns are consistent with a story in which a narrower network leads patients who expect to get care from the excluded providers to switch plans to keep access to their hospital or doctor. Consistent with this interpretation, 91% of the 2012 switchers (and 98% of Partners patient switchers) shift to one of the two plans that still covers Partners (CeltiCare and NHP). This story is quite intuitive.

[26]Switching out rates were also high in 2009, reflecting unusually large increases in Network Health's enrollee premiums from 2008-09. Figure 3 shows that in contrast with 2012, switching in 2009 was not correlated with proxies of demand for Partners hospitals.

[27]Risk adjustment was first implemented in 2010, so it is only available for prior-year costs starting in 2011.

[28]Another way of viewing these patterns is to flip the conditional probabilities and ask what share of switchers each group represents. Partners patients represent 18% of Network Health enrollees in 2011 but (because they are so much more likely to leave) comprise 67% of switchers out. Other dropped hospitals' patients represent 8% of 2011 enrollees but 17% of switchers out. Thus, these two groups together comprise the vast majority (84%) of switchers out in 2012.

Figure 2: Plan Switching and Selection for Network Health (around 2012 network change)

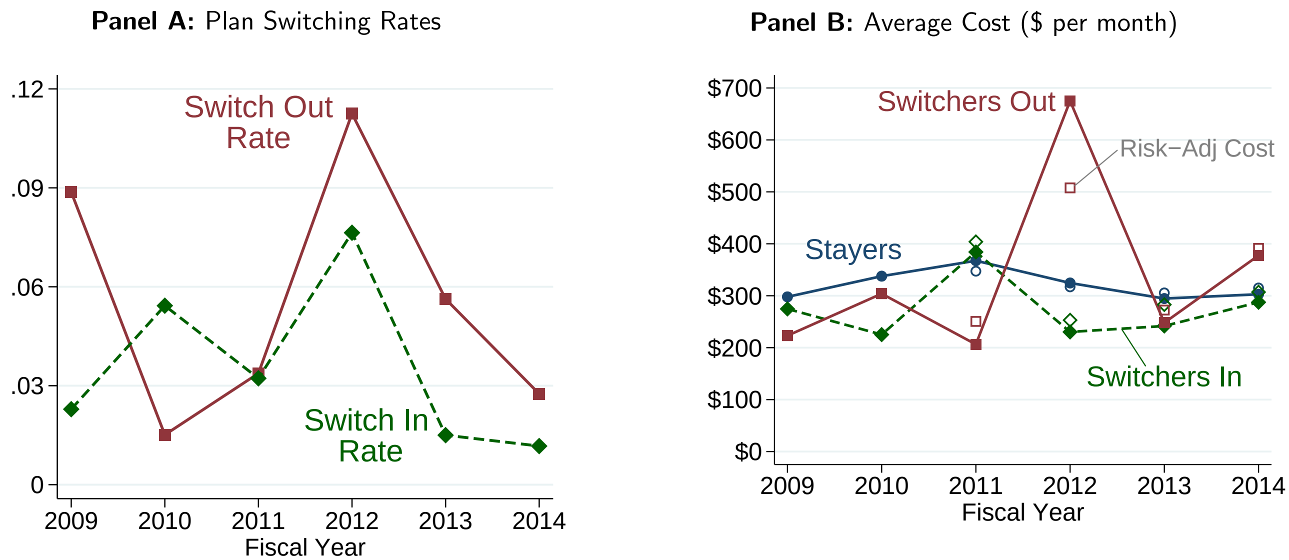

NOTE: These figures show switching and selection patterns for Network Health over time and especially around its 2012 network narrowing. Panel A shows the rate of switching in and out of Network Health at each year's open enrollment. These rates are defined as the number of switchers in/out divided by the same denominator – the number of continuous market enrollees in Network Health at the end of the prior year – so their levels are comparable. Panel B shows the average prior-year cost of stayers, switchers out of, and switchers into Network Health. The connected series (with solid points) are raw average costs. The open points are risk-adjusted costs using the exchange's method: a group's average costs divided by a group's average risk score. Risk adjustment was implemented starting in 2010, so it is only available for prior-year costs starting in 2011.

The fact that it holds for patients both of Partners and the other dropped hospitals suggests a general mechanism, not something specific to star hospitals. This occurs despite the well-known fact of low switching rates in health insurance settings (e.g., Handel, 2013; Ericson, 2014). One possible reason – for which there is anecdotal evidence from my discussions with providers – is that the dropped hospitals contact their patients and encourage them to switch plans. This provision of advice by providers may represent an important mechanism through which plan networks influence enrollee choices.

The final logical step in connecting provider demand to adverse selection is to show that these groups tend to have high risk-adjusted costs. Appendix Table 13 shows evidence of this. Among all continuing 2011 Network Health enrollees (switchers plus stayers), both raw and risk-adjusted costs are higher for the groups most likely to switch out – people living nearby Partners and patients of Partners or the other dropped hospitals. The highest-cost group are Partners patients, with risk-adjusted costs of $564 per month, or 63% above average. Of course, this analysis does not explain *why* the switching groups had high costs, a question that matters for interpreting the findings. I return to this issue in Section 5.

4.2 Robustness Analyses on Adverse Selection Findings

The evidence so far has focused on plan switching patterns for Network Health's current enrollees at the end of 2011. This section implements three analyses to check the robustness of these findings:

Figure 3: Plan Switching Out Rates for Network Health, by Year

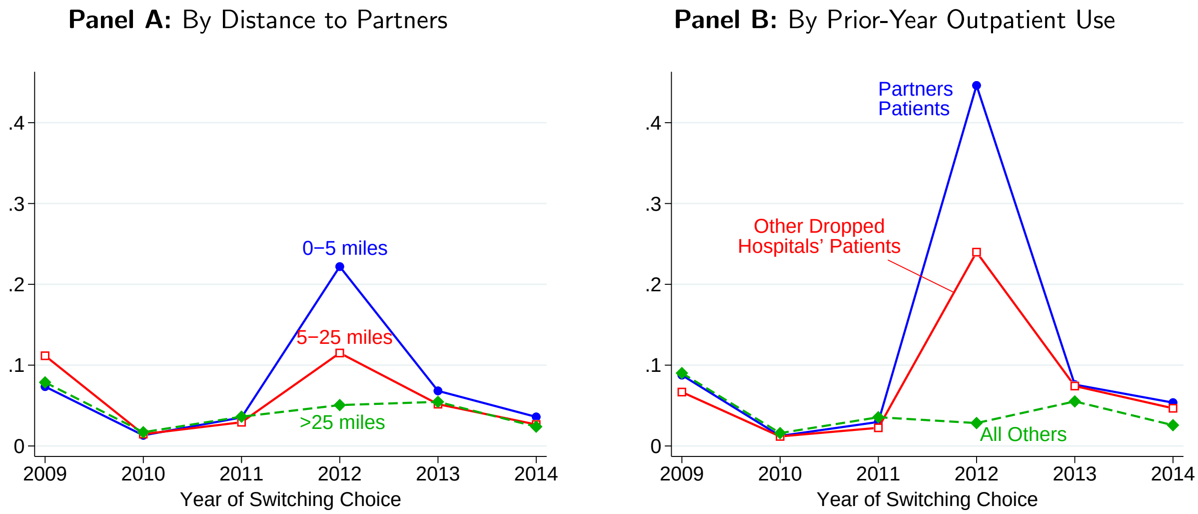

NOTE: These figures show switching out rates for enrollees of Network Health at each year's open enrollment, separately by variables likely to correlate with demand for the providers dropped from network in 2012. Panel A shows switching rates by enrollee distance to the nearest Partners hospital. Panel B shows switching rates by whether enrollees used Partners or another dropped hospital for (non-emergency room) outpatient care in the prior year. In both cases, the groups likely to have higher demand for the dropped hospitals are more likely to switch out of Network Health in 2012, and this represents a clear departure from switching patterns in other years.

(1) examining new enrollee choices, which are not subject to inertia; (2) studying switching by zero-premium enrollees, for whom there is no concurrent change in Network Health's premium that could affect results; and (3) showing similar evidence from CeltiCare's 2014 exclusion of Partners from its network.

Evidence from New Enrollee Choices While switching behavior provides the cleanest evidence of adverse selection, another important channel is changing plan demand among "new enrollees" entering the exchange. I briefly provide evidence of similar selection patterns among this group; their choices also enter the plan demand estimates in the structural model. A challenge with studying new enrollees is that, because they newly join the market, I often lack data on their costs and provider use *prior to* the network change (and outcomes after the change could be directly influenced by it). Therefore, when I study cost/utilization variables, I restrict to the subset of "re-enrollees" who have a prior CommCare enrollment spell that ended before 2012. I use this prior spell to measure provider use and costs. In addition, because of the 2012 limited choice policy for below-poverty new/re-enrollees (see Section 3.3), I limit the analysis to above-poverty enrollees who have unrestricted choice.

Figure 4 shows evidence of changing demand for Network Health in 2012 that is correlated with markers of provider demand – just as in the switching findings in Figure 3. Each point on the graphs represents Network Health's market share for the group of new enrollees joining the exchange in a given bimonthly period. Panel A breaks out market shares by enrollee distance to the nearest Partners hospital. While demand increases in 2012 for all groups – reflecting the plan's premium decrease – the jump is much smaller for people living within 5 miles of a Partners hospital. Panel B shows even

Figure 4: Network Health's New Enrollee Market Share around 2012 Change

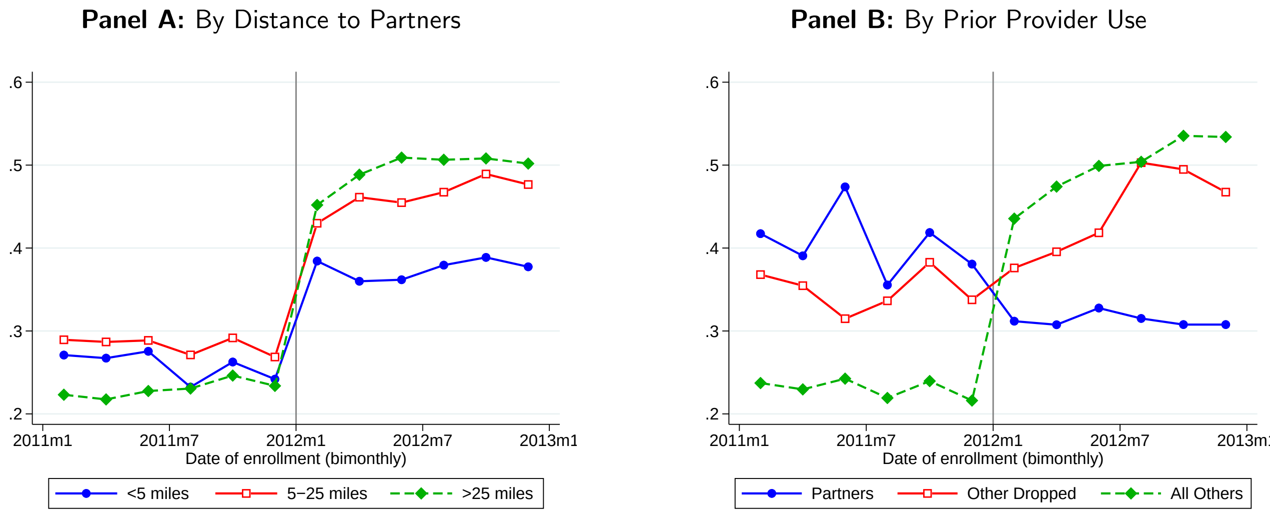

NOTE: These figures show evidence of changes in new enrollees' demand for Network Health in 2012 that are correlated with valuation for the Partners and other dropped hospitals. Each point on the figures is the market share who choose Network Health among above-poverty new enrollees joining the exchange in a given (bimonthly) period. The sample is restricted to above-poverty enrollees who are not subject to the 2012+ limited choice policy. Panel A divides enrollees by proximity to the nearest Partners hospital. Panel B divides enrollees by use of the dropped hospitals during a prior enrollment spell, with the sample limited to re-enrollees with a previous spell. In both panels, market shares increase in 2012 for groups least likely to value the dropped hospitals (reflecting Network Health's premium decrease) but increase much less or decline for groups more likely to value the hospitals.

starker results breaking out demand among re-enrollees based on use of the dropped hospitals during their prior spell. While market shares for the "all others" group (who did not use Partners or another dropped hospital) more than doubles from ˜25% in 2011 to over 50% in late 2012, shares for Partners patients decline in 2012. Shares for other dropped hospitals' patients increase but by much less than for the "all others" group.

These results show that the impact of the network change on plan demand was not limited to plan switching but also had a major effect on new enrollee choices. Appendix Figure 13 shows that these demand shifts were correlated with proxies for costs in a way suggesting more favorable selection. Following the change, the plan's new enrollees' average risk score falls and its re-enrollees' prior-spell average cost decreases – implying that older and higher-cost enrollees select away from the plan. Although this evidence is more limited than for switching, it again is consistent with the basic adverse selection story.

Plan Switching for Zero-Premium Enrollees The selection changes for Network Health in 2012 reflect a combination of its narrower network and lower premium, which are part of the same strategic bundle. However, a natural question is whether the results are *entirely* driven by the lower premium, rather than the network shift. The CommCare setting provides an easy way to test this by examining switching patterns for below-poverty enrollees for whom all plans are free (both before and after 2012). Importantly, existing below-poverty enrollees were not subject to the limited choice policy (which applied only to new enrollees) so could switch freely.

Appendix Figures 17-18 replicate Figures 2-3 with the sample limited to below-poverty enrollees. Both switching out and cost patterns for stayers/switchers out are quite similar to the full sample. The one meaningful difference is instructive: there is no spike in low-cost below-poverty enrollees *switching into* Network Health in 2012, consistent with the lack of a premium incentive to do so. This suggests that the network and premium changes work together in driving selection incentives: the narrower network pushes out high-cost enrollees who care about provider choice, while the lower premium pulls in low-cost enrollees who are price-sensitive. These findings suggest that adverse selection on networks is likely relevant in settings without premiums (e.g., Medicaid managed care) but may be more muted.

Evidence from CeltiCare 2014 Dropping of Partners The analysis so far relies on a single network change for Network Health in 2012. It is reasonable to ask whether this is a fluke. To provide evidence, I examine the only other CommCare network change involving the star Partners system: when CeltiCare drops Partners at the start of fiscal year 2014. This change is at the tail end of my data period, limiting the analyses I can do (e.g., the claims data for 2014 are incomplete). Nonetheless, to provide an additional source of evidence, I replicate the analyses of Figures 2-4 above for CeltiCare.

The results are shown in Appendix E.2 (see Figures 14-16). All of the main selection findings carry over to CeltiCare in 2014. Specifically: (1) CeltiCare experiences a high switching out rate in 2014, with switchers out having high raw and risk-adjusted costs; (2) switching rates are strongly correlated with proximity to Partners and prior-year use of Partners, and (3) CeltiCare's demand among new enrollees shows similar patterns (falling for Partners patients and people living nearby a Partners hospital, while rising for others). Together, these results suggest that Network Health's 2012 experience was not an idiosyncratic event but representative of generalizable patterns of selection based on star hospital coverage.

4.3 Understanding Demand for Star Providers and the Role of State Dependence

Why do some individuals exhibit high demand for the star providers, as exhibited in their willingness to switch plans to retain access? What role do state dependence and heterogeneity play? This issue is relevant for interpreting the short- vs. long-term patient welfare losses from the narrower networks. While the data do not provide a good way to precisely decompose the precise contribution of each channel, this section presents evidence suggesting that both are involved.

Start by noting that the fact that people switch plans does *not* distinguish state dependence from heterogeneity. While switching out of Network Health in 2012 – which involves an administrative hassle and often paying a higher premium[29] – suggests a desire to keep one's hospital/doctor, there are two reasons people may have this preference. First, they may be "matched" to their provider based on *persistent heterogeneity* in factors that make the provider more attractive: good care for their

[29] Below-poverty enrollees could switch to any plan and still pay zero premium, but above-poverty enrollees faced a choice of two plans that covered Partners: (1) NHP, whose premium was $21-51 per month higher than Network Health (depending on income), or (2) CeltiCare, which cost the same as Network Health but had a much narrower network in other ways (see Appendix Figure 11) and a worse reputation (as indicated in the plan demand estimates in Table 10). Interestingly, switching rates for below- and above-poverty enrollees were quite similar.

condition, greater convenience, or other factors. Alternatively, they may simply not want to switch providers, especially if they have a good relationship or are in the middle of an active treatment regime. These explanations are examples of state dependence because they arise from *past treatment history*. Notice that they may be still be quite important to patients and even clinically meaningful in the sense that breaking the relationship harms a patient's health (see Sabety, 2020). But their key feature is that they are rooted in past history that might have been different and whose importance may fade over time.

To examine these mechanisms, I dig deeper into who switches plans in response to Network Health's 2012 network change. As in Section 4.1, this section limits the sample to current Network Health enrollees at the end of 2011 and runs regressions to analyze who switches out of the plan at the start of 2012.

Evidence of Heterogeneity Table 2 shows (binary) logit regressions, with the outcome variable in columns (1)-(2) an indicator for switching out of Network Health. The x-variables are various characteristics that may predict heterogeneous value for the Partners hospitals or other dropped providers: distance (i.e., convenience), medical conditions, and demographics. To aid interpretation, I report odds ratios (which equal e^{β} of the underlying logit coefficients, β).

Column (1) shows results *without* controlling for prior provider use. This model therefore sheds light on whether there is "matching" on characteristics associated with provider demand in a history-unconditional sense. The estimates indicate strong evidence of this matching. One clear factor is convenience: individuals are more likely to switch out if they live closer to a Partners hospital or another dropped hospital, with odds >7x higher for people living within 2 miles and gradually declining with further distance. A second set of factors are medical risk and conditions. These matter because the star hospitals are known for their advanced care for the sickest patients – the explicit criteria on which the *U.S. News* rankings are based. Switching rises with age (consistent with age as a risk factor) and with observed medical conditions. Having any chronic or acute illness increases switching odds by 68% and 42%, respectively. On top of these, there are sizable further effects of having a risk score in the top 5% (+45%) and having cancer (+110%). Cancer is notable because Brigham & Women's Hospital is clinically integrated with Dana Farber Cancer Institute, the region's top cancer hospital, making it difficult to get care at Dana Farber without access to Brigham's facilities.

These differences imply that in an unconditional sense, provider preferences revealed in plan switching reflect real heterogeneity in value for the star hospitals. However, it is important to interpret these findings with care. While they indicate that there *is* real sorting on persistent determinants of provider demand (i.e., heterogeneity), they do *not* rule out state dependence – or even suggest that it is unimportant. It is a mistake to think of this as an "either/or" story; rather a "both/and" approach is more appropriate. Indeed, heterogeneity and state dependence are likely deeply intertwined. Individuals may *initially* sort into becoming a Partners patient based on real heterogeneity (e.g., convenience or sickness) but remain loyal to Partners because of a mix of heterogeneity and state dependence (e.g., a switching cost or the relationship's value). Columns (2)-(3) of Table 2 indicate support for both

Table 2: Heterogeneity in Likelihood to Switch Out after 2012 Network Narrowing

Variable	Outcome: Switch Out of Network Health — Unconditional (1) Odds Ratio	(S.E.)	Controlling for Patient Status (2) Odds Ratio	(S.E.)	Outcome: Being a Partners Patient (3) Odds Ratio	(S.E.)
Distance to Partners Hospital						
0-2 miles	7.24	(0.45)**	2.17	(0.16)**	40.61	(2.94)**
2-5 miles	4.83	(0.24)**	1.96	(0.12)**	22.93	(1.44)**
5-10 miles	2.68	(0.15)**	1.29	(0.08)**	13.45	(0.86)**
10-20 miles	2.40	(0.15)**	1.25	(0.08)**	8.53	(0.59)**
20-30 miles	1.25	(0.08)**	1.09	(0.07)	3.14	(0.23)**
> 30 miles	*(omitted = 1.0)*		*(omitted = 1.0)*		*(omitted = 1.0)*	
Medical Risk and Conditions (during 2011)						
Age (years/10)	1.21	(0.02)**	1.23	(0.02)**	1.04	(0.01)**
Any Chronic Illness	1.68	(0.07)**	1.01	(0.05)	2.26	(0.09)**
Any Acute Illness	1.42	(0.06)**	0.64	(0.03)**	3.34	(0.16)**
Risk Score in top 5%	1.45	(0.10)**	1.17	(0.09)*	1.59	(0.10)**
Cancer	2.10	(0.17)**	1.64	(0.15)**	2.56	(0.21)**
Cardiovascular	1.51	(0.12)**	1.26	(0.11)**	1.55	(0.12)**
Diabetes	1.05	(0.06)	1.08	(0.07)	0.95	(0.05)
Lung Disease	1.18	(0.08)*	1.07	(0.08)	1.19	(0.08)**
Mental Health	1.04	(0.06)	1.04	(0.06)	1.08	(0.05)
Pregnancy	0.63	(0.19)	0.46	(0.15)*	1.53	(0.33)*
Patient at Dropped Providers during 2011						
Partners Provider	---		23.25	(1.14)**	---	
Other Dropped Provider	---		12.24	(0.71)**	---	
Observations	41,918		41,918		41,918	
Pseudo-R^2	0.105		0.305		0.232	

* Statistical difference from an odds ratio of 1.0 is indicated with ** (1% level) and * (5% level).

NOTE: The table reports estimates of binary logit regressions for the outcome of switching out of Network Health in 2012 (columns 1-2) and being a Partners patient for outpatient care in 2011 (column 3). The sample consists of current enrollees in Network Health as of the end of 2011 who choose whether or not to switch plans at the start of 2012. The table reports logit odds ratios, equal to e^{β} of the underlying logit coefficients β. Distance is defined as driving distance to the closest Partners hospital. All medical conditions are defined based on diagnoses on 2011 claims. Any chronic and acute illnesses are defined based on a categorization shared with me by Kaushik Ghosh and David Cutler. The specific illnesses are based on a categorization of diagnoses entering the HCC risk score model. The top 5% risk score category is based on CommCare's risk score as calculated from 2011 claims data. In addition to the variables shown above, the model includes controls for gender and income group.

stories. Column (3) reports a logit for the outcome of being a Partners patient in 2011 and finds that there is strong sorting based on convenience and medical conditions. Column (2) shows that even after controlling for being a Partners patient in 2011 – which is by far the strongest predictor of switching, with an odds ratio of 23.25 – convenience still predicts switching. Age, high risk score, cancer, and cardiovascular disease also predict higher switching. But interestingly, acute illness and pregnancy during 2011 have odds ratios significantly below one (0.64 and 0.46), indicating these groups are less likely to switch (conditional on other covariates). This suggests forward looking behavior as individuals care less about provider access once they have recovered from temporary conditions.

Overall, this evidence is most consistent with a role for *both* heterogeneity and state dependence. Importantly, this suggests that patients likely suffer real utility losses both in the short and long run if they lose access to their preferred providers. Someone who has cancer or lives nearby a Partners hospital loses out from the narrower network, even after they switch to a new provider. As long as provider sorting is partly based on persistent factors (either initially or dynamically), there are long-run welfare implications. Of course, state dependence also matters because it *amplifies* how much patients care today about keeping their doctor, relative to the long run.

Evidence of State Dependence The findings so far are suggestive that state dependence is relevant. To provide stronger evidence, I examine the role of of a more detailed treatment history variable: the *recency* of the latest visit to a physician of Partners or another dropped provider. The model I have in mind is one where a patient's loyalty is determined by the strength of the patient-doctor relationship. That relationship, in turn, is strongest when recently renewed through an in-person office visit and decays gradually as time elapses without an interaction. Of course, the main concern in testing this story is that visit recency correlates with illness – sicker people get care more frequently – so I will do my best to control for sickness in the analysis.

Figure 5 shows how probability of switching out of Network Health at the start of 2012 varies with months elapsed since the patient's last office visit to Partners or another dropped hospital's physician. The sample is split among Partners patients (blue), patients of other providers dropped by Network Health in 2012 (red), and as a control group, patients of all other providers who are not dropped (green).[30] The plot shows binned predicted probabilities from logit regressions (separately by patient group) after controlling for a detailed set of demographic, health status, and distance-to-provider variables (see figure notes), along with quadratic best-fit curves. Appendix Table 14 reports the numerical estimates and shows robustness to the controls included.

For patients of Partners or another dropped provider, there is a steep relationship between visit recency and the likelihood of switching out of Network Health in 2012. Among patients who visited Partners in the past 1-2 months, 62-71% switch plans – an extremely high rate for insurance choice where inertia is the norm. This declines to 52-56% for patients with a visit 3-6 months prior, 43-

[30] The analysis excludes about 19% of individuals do not have any observed physician visits prior to the start of 2012. Among the remaining sample, 13% have a prior Partners visit and 4% have a prior visit to another dropped hospital's physician, with a small number of overlaps (0.3%) classified as Partners patients. The x-variable is defined as months since the last visit to the provider in the indicated system (Partners or other dropped) – i.e., it does not count more recent visits to other providers.

Figure 5: Switching Rate Out of Network Health, by Recency of Last Provider Visit

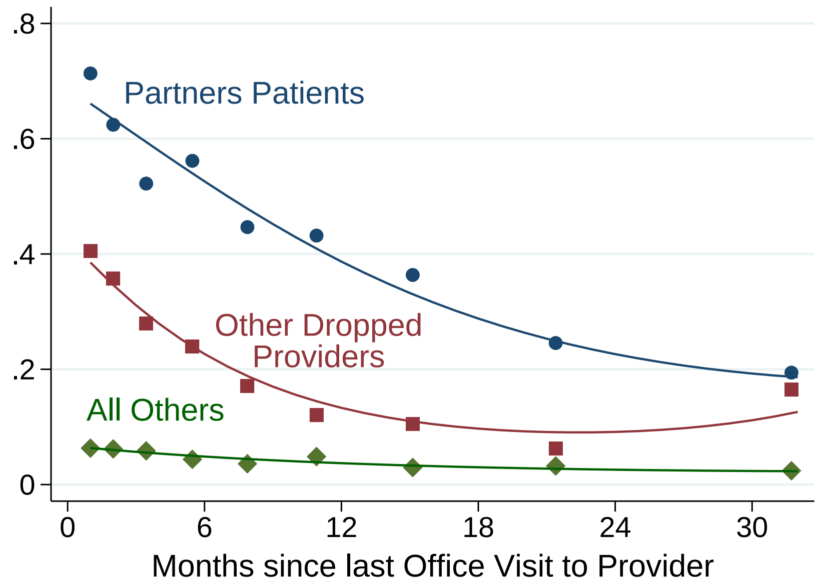

NOTE: The plot shows how plan switching rates out of Network Health in 2012 relate to the recency of a physician office visit with the indicated provider. Individuals are categorized into Partners patients (blue circles), patients of another dropped hospital (red squares), and all other patients (green diamonds) based on prior physician office visits in the claims data. Individuals with no prior office visits in the data are excluded, and a small number (0.3%) of overlaps between Partners and other dropped providers' patients are classified as Partners patients. The x-axis is recency (as of the start of 2012) of the latest physician office visit to the indicated provider (e.g., Partners for the Partners patients). The numbers shown are predicted probabilities for recency bins from logit regressions, controlling for demographics (age, gender, income group), medical risk variables (chronic condition dummies and vigintiles HCC risk score), and distance to Partners and other dropped hospitals. Separate regressions are run for each patient group, and predicted probabilities are evaluated at the mean of control variables. The lines are quadratic best-fit curves.

45% for patients with a visit 7-12 months prior, and gradually down to 19% for patients whose most recent visit is 25+ months prior (the final plotted bin). There is a similar pattern for patients of other dropped providers, albeit at a lower level of switching. For all other patients, switching is only modestly related to visit recency.

These results in Figure 5 suggest that consumers' willingness to switch plans to keep their provider is influenced not just by the existence of a relationship but by how recently it has been renewed. They are strongly consistent with history (i.e., state dependence) mattering for provider preferences, and particularly so for the star hospitals. While not perfect evidence – visit recency is not randomly assigned – the patterns are difficult to explain with other stories. The results control for detailed medical risk variables (along with demographics and distance), suggesting that recency is not merely proxying for sickness. Results are also not sensitive to which controls are included (see Appendix Table 14). Moreover, the patterns are only present based on recency of visits to the dropped providers, not to other providers. Thus, the most likely explanation is that past experience with a provider matters – and matters more so when that experience is recent.

5 Understanding Costs Driving Adverse Selection

The evidence so far indicates that insurers have a selection incentive to narrow their networks by excluding high-cost star providers. In addition to any direct cost reductions, the narrower network drives away consumers who want access to the excluded providers – especially existing patients who live nearby or have severe conditions – and these individuals tend to have been high-cost and unprofitable even with risk adjustment. This raises the natural question of *why* these individuals had high costs? The model in Section 2 suggested two (non-exclusive) channels: (1) they may be unobservably sicker patients, and/or (2) they may have differentially large cost increases (moral hazard) when given access to the star providers. Cost increases, in turn, may arise via the star providers receiving higher prices or delivering higher quantity of care.

Because these channels have different implications, this section provides additional evidence on their relevance. Section 5.1 begins with a method to decompose spending into prices vs. quantities. Section 5.2 shows evidence on the role of sickness and provider prices. Section 5.3 shows evidence of moral hazard and analyzes the role of price vs. quantity.

5.1 Price-Quantity Decomposition

To understand the sources of cost variation, I decompose medical spending into prices vs. quantities. I focus on inpatient and outpatient care for which I can clearly observe the unit of service and payment per service. This "decomposition sample" comprises the vast majority of hospital-based care and about two-thirds of overall medical costs, with the main excluded cost being prescription drugs.[31]

Spending decompositions can be done in a variety of ways. My approach is to define quantity as "price-standardized" utilization, and price as the remaining factor accounting for observed spending. To see how this works, note that observed insurer spending (C_{it}) for individual i in year t equals the sum of payments for various distinct medical services, $s \in \{1, ..., S\}$, whose definition I discuss below. Letting $a_{it} \in A_{it}$ index each instance of i receiving care in year t, $C_{it} = \sum_{a_{it} \in A_{it}} Paid_{a_{it},s(a_{it})}$. Each payment can in turn be broken down into the "quantity" (Q_s) associated with service s and the remaining "price" factor ($P_{a_{it}s}$): $Paid_{a_{it},s} = Q_{s(a_{it})} P_{a_{it},s(a_{it})}$. While there is no perfect way to define quantity, a natural method is to use the mean payment for a service (across all insurers and years) so that the remaining price measure is centered around 1.0 for each service. With this definition of Q_s, total quantity of care used by person i in year t equals:

$$Q_{i,t} = \sum_{a_{it} \in A_{it}} Q_{s(a_{it})} \tag{7}$$

[31]See Appendix C for the details of the sample construction and decomposition method. I exclude drugs because their prices should not be related to the hospital network and because of the challenge of observing true prices due to unobserved "rebates" from pharmaceutical companies to insurers. In addition to drugs, the sample omits inpatient care in specialty hospitals (e.g., psychiatric hospitals and residential facilities) and outpatient care paid via a method besides fee-for-service.

and the individual's average price of care is:

$$P_{i,t} \equiv \frac{C_{i,t}}{Q_{i,t}} = \frac{\sum_{a_{it}} \left[Q_{s(a_{it})} P_{a_{it},s(a_{it})} \right]}{\sum_{a_{it}} Q_{s(a_{it})}} \tag{8}$$

which is a quantity-weighted average of prices across all services received among individuals with positive quantity.[32] To make the units of price and quantity comparable, I can divide by its mean in a sample ($\overline{Q}$) to calculate relative quantity $Q^{rel}_{i,t} \equiv Q_{i,t}/\overline{Q}$. The final cost decomposition is:

$$C_{i,t} = Q_{i,t} \times P_{i,t} = \overline{Q} \times Q^{rel}_{i,t} \times P_{i,t} \tag{9}$$

A key step in this decomposition is defining the unit of medical services, s. I do so slightly differently for outpatient and inpatient care. Appendix C reports the details of the method. To summarize, for outpatient care, I define services based on procedure codes (HCPCS codes, as defined by Medicare), the standard measure used in previous work (e.g., Clemens and Gottlieb, 2017; Brot-Goldberg et al., 2017). I further interact these codes with the type of bill/provider to allow quantity to vary across settings (facility vs. non-facility) and type of care (e.g., medical vs. behavioral health vs. dental care). For inpatient care, the service unit is an admission for a particular diagnosis-related group (DRG) or diagnosis (if DRG is not used for payment), adjusted for patient severity observables. In practice, I implement this definition via a regression model, following a method similar to past work (Cooper et al., 2019).

5.2 Understanding High Costs: Sickness and Use of High-Price Providers

I now apply the decomposition just described to shed light on the reasons for the high costs of switchers who leave Network Health in 2012. Table 3 compares stayers versus switchers on various cost measures, with columns (1)-(3) comparing group averages and columns (4)-(5) showing how differences are affected by controlling for medical risk to indicate the role of sickness in explaining gaps.

Differences on Medical Risk (Observed and Unobserved) The first row of Table 3 shows overall switcher-stayer cost differences (as in Figure 2B). Before controlling for medical risk, column (3) shows that switchers are \$353 per month (or 107%) more expensive than stayers. A part of this difference reflects observed medical risk, based on CommCare's risk score used for risk adjustment. Controlling for this measure shrinks the gap to \$286 per month (or 87%). This is consistent with the findings in Table 2 that switchers were more likely to have chronic medical conditions that are key inputs to risk adjustment.

Does the remaining gap reflect *unobserved* medical risk? A simple way to test for this is to examine a richer measure of medical risk. The CommCare risk score is retrospective (based only on lagged

[32] This price index is analogous to a standard Paasche price index, treating the "base-period" price as Q_s. Notice that $P_{i,t}$ can only be measured for individuals with positive quantity; all price variation results are conditional on this sample (about 77% of enrollee-years for outpatient and overall costs, though only 4% for inpatient care). When calculating average price for a group of people, I weight by individual quantities so that the product of average quantity and price equals average cost.

diagnosis information) so misses diagnoses observable only in current-year data. The "concurrent" HCC risk score used by the ACA, however, captures this component of risk (see discussion in Section 3.1). Column (5) shows that controlling for the HCC score further shrinks the gap, indicating that switchers are also sicker on factors not captured by CommCare's risk adjustment. But a substantial gap of $186 (or 56%) remains. The top rows of Panels B-C show that patterns are similar for inpatient and outpatient care, the two components for which I can decompose costs. Switchers are more costly than stayers – with a particularly large 220% gap for inpatient care – but gaps shrink as better medical risk controls are added. However, the remaining difference in column (5) is still 78% for inpatient and 70% for outpatient care.

Cost Decomposition: Quantity versus Price Does the remaining cost gap after controlling for medical risk represent higher quantity of care or higher prices for similar care? Higher price would clearly reflect the novel channel highlighted in my theory. Quantity, however, would be more ambiguous, reflecting either further unobserved sickness or providers' causal impacts on treatment intensity.

The second rows of Panels B-C of Table 3 show the role of quantity differences, relative to the mean for stayers normalized to 1.0 ($\overline{Q} = \overline{Q}_{stayer}$). For both types of care, quantity differences account for a large portion of raw cost gaps. For instance, switchers have higher inpatient quantity by 1.49 (indicating 149% greater use of care), or about two-thirds of the total inpatient cost difference. However, higher raw quantity may simply indicate that stayers are sicker. Columns (4)-(5) show that the gaps shrink substantially after controlling for medical risk, down to 32% for inpatient and 64% for outpatient care in column (5). As noted above, these remaining quantity gaps could reflect either further unobserved risk or provider effects on treatment intensity.

The final section of Panels B-C shows price differences. Here the patterns differ for inpatient versus outpatient care. For inpatient care (Panel B), switchers use higher-price hospitals, accounting for a price factor difference of 0.29 (or 28% higher than stayers). This is driven by the switchers' >4x larger share of choosing Partners hospitals (69% vs. 15%), whose average price for switchers is 1.45 (or 45% above the mean). Moreover, both the price gap and the Partners share are nearly unaffected by medical risk controls in columns (4)-(5). This is consistent with a key idea in the model: that risk adjustment is unlikely to offset cost differences arising from different provider prices and choices, which are not purely driven by medical status. Thus, although prices account for only a small portion of *raw* cost differences, they account for almost half of *risk-adjusted* differences in column (5).

The patterns for outpatient care are somewhat different. Switchers continue to choose Partners providers for a much larger share of care (33% vs. 6%), a >5x difference. But Partners' outpatient care prices are not high (0.97, or 3% *below* average), so average outpatient prices for switchers are quite similar to stayers. Instead, nearly all of switchers' higher outpatient spending comes from higher quantity. Moreover risk adjustment narrows quantity gaps by much less (proportionally) than it did for inpatient care. This could indicate either that outpatient care is more affected by hard-to-measure risk factors or that Partners has a causal effect on outpatient care through the intensity of services provided. I explore these causal effects further in the next subsection.

Summary Overall, these results are consistent with a nuanced view of the source of high costs that nonetheless indicate the power of the hospital network as a driver of adverse selection. Dropping the star hospitals reduces demand among consumers who are high costs along *multiple dimensions of costs*: observable risk (as captured by risk scores), unobserved risk (as captured by HCC risk scores), high quantity of care conditional on risk, and greater use of high-price hospitals for inpatient care. Better measures of medical risk (e.g., the concurrent HCC score vs. the lagged CommCare score) shift a larger share of costs into the "observed risk" category, which has an important effect on mitigating selection incentives. But they miss a substantial remaining portion of costs coming from both higher quantities and prices of care delivered by the excluded medical providers. These results indicates the durability of selection incentives associated with covering/excluding star hospitals.

5.3 Causal Cost Changes (Moral Hazard)

My model in Section 2 emphasizes a particular channel for adverse selection: selection by people with *high incremental costs due to star hospital coverage* (i.e., high moral hazard), creating a form of "selection on moral hazard." To test this prediction, I now examine whether dropping Partners had a *causal effect* of reducing enrollee-level medical spending and how that causal effect varies across enrollees. In addition to testing this idea, these estimates form a key part of the structural cost model presented in Section 6.2.

To do so, I again draw on the natural experiment of Network Health's 2012 network narrowing. Instead of studying plan switching, I examine cost changes for the panel of "stayers" continuously enrolled in Network Health from 2011-2012, relative to a control group of stayers in other plans (whose networks were stable). Limiting the sample to stayers and the 2011-12 period, I setup a Poisson regression with individual and time fixed effects. The estimating equation is:

$$E\left(C_{i,j,t}\right) = \exp\left(\alpha_i + \beta_t\left(Z_i\right) + \gamma\left(Z_i\right) \cdot 1\left\{j = NH, t \geq 2012\right\}\right) \tag{10}$$

where $C_{i,j,t}$ is insurer cost on individual i at time t, α_i is an enrollee fixed effect (which is divided out and not estimated), $\beta_t(.)$ are time fixed effects that capture trends for the control group, and Z_i are enrollee characteristics on which time trends and causal effects may vary. Regression (10) is estimated by maximum likelihood (using "xtpoisson, fe" in Stata), with cluster-robust standard errors at the i level. The coefficients of interest are $\gamma\left(Z_i\right)$, which capture the differential cost change for Network Health stayers in 2012. Note that (10) is analogous to standard difference-in-differences but in a non-linear model.[33] The implied (multiplicative) effect on costs equals $dC_i = \exp\left(\hat{\gamma}\left(Z_i\right)\right)$, and the percent change is $dC_i - 1$. I also estimate event study versions of (10) that allow $\gamma_t(.)$ to vary with time.

[33]I adopt a Poisson specification since it is natural to think that networks affect costs proportionally to an individual's baseline spending and also to aid decomposing effects into price vs. quantity. However, all main results are robust to using a linear fixed effects specification.

Table 3: Decomposition of Switchers' High Costs: Quantity vs. Price

Outcome	Average Value: Stayers	Average Value: Switchers Out	Difference: Switchers Out - Stayers: Raw Diff.	With Risk Controls: CommCare	With Risk Controls: HCC Score
	(1)	(2)	(3)	(4)	(5)
Panel A: Overall Spending					
Total Spending	\$330.4	\$683.8	\$353.4**	\$286.1**	\$185.6**
			(27.3)	(25.1)	(16.0)
Panel B: Inpatient Care					
IP Cost (\$ per month)	\$47.7	\$152.8	\$105.1**	\$87.9**	\$37.0**
			(15.8)	(13.9)	(8.9)
Quantity (relative)	1.00	2.49	1.49**	1.19**	0.32*
			(0.23)	(0.21)	(0.14)
Price of Care	1.02	1.31	0.29**	0.28**	0.25**
			(0.03)	(0.03)	(0.03)
Share at Partners	*15%*	*69%*	*0.53***	*0.53***	*0.50***
			(0.03)	*(0.03)*	*(0.03)*
Price at Partners	*1.33*	*1.45*			
Price at All Others	*0.96*	*0.99*			
Panel C: Outpatient Care					
OP Cost (\$ per month)	\$153.3	\$301.1	\$147.8**	\$125.0**	\$108.0**
			(8.3)	(8.0)	(6.8)
Quantity (relative)	1.00	1.92	0.92**	0.77**	0.64**
			(0.05)	(0.05)	(0.05)
Price of Care	0.95	0.97	0.024*	0.023*	0.031**
			(0.011)	(0.011)	(0.011)
Share at Partners	*6%*	*33%*	*0.27***	*0.27***	*0.27***
			(0.01)	*(0.01)*	*(0.01)*
Price at Partners	*1.04*	*0.97*			
Price at All Others	*0.94*	*0.97*			
Number of Enrollees	37,201	4,716			

NOTE: Statistical difference from zero is indicated with ** (1% level) and * (5% level). The table shows a decomposition of costs differences between stayers and switchers out of Network Health in 2012 (groups defined in the text) following the price-quantity decomposition in Section 5.1. Columns (1)-(2) report group averages; column (3) shows the raw differences between groups (with standard errors in parentheses), and columns (4)-(5) show differences after controlling for the CommCare and HCC risk scores (see definitions in text). Panel A shows total medical spending. Panel B shows inpatient care, with total spending and a decomposition into relative quantity and price. Panel C shows the same for outpatient care. The sum of inpatient and outpatient care is less than total spending because of other costs not covered by the decomposition.

Figure 6: Cost Reductions for Stayers after 2012 Network Change

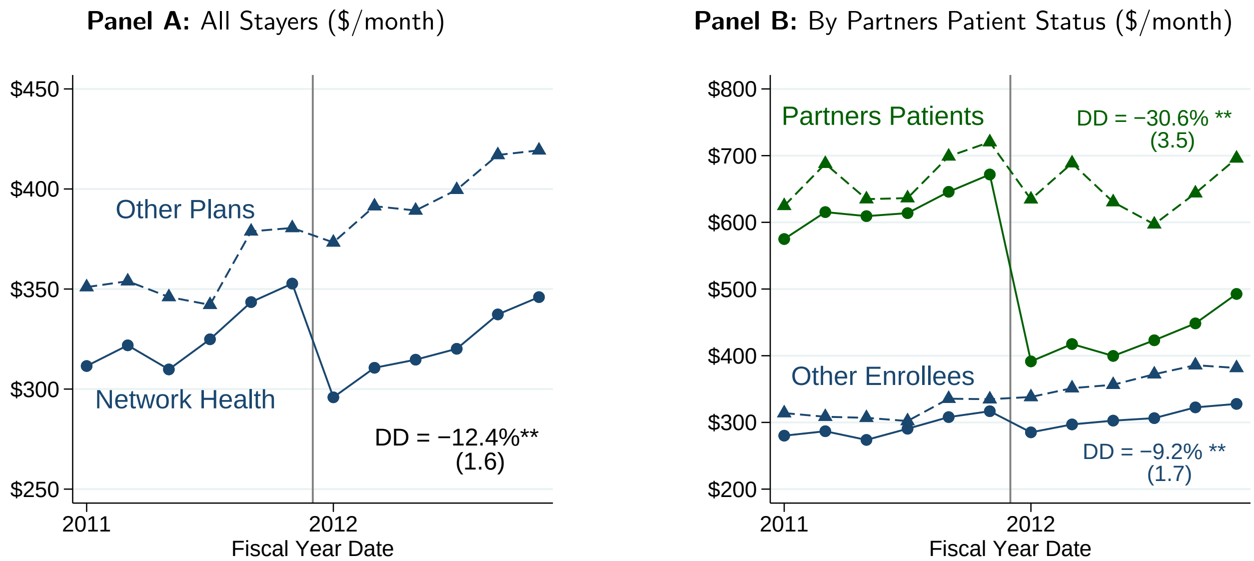

NOTE: These graphs show estimates from cost regressions with individual fixed effects corresponding to the event study version of equation (10). The sample is "stayers" continuously enrolled in Network Health or other plans between 2011 and 2012, when Network Health narrows its network. The outcome variable is insurer costs (in $ per month) averaged over bimonthly periods. The graphed points correspond to estimates of $\exp(\overline{\alpha}_{Oth} + \beta_t)$ (for other plans) and $\exp(\overline{\alpha}_{NH} + \beta_t + \gamma_t)$ (for Network Health). I also report the DD estimate of the percent change in costs ($= \exp(\gamma) - 1$) and its standard error. Standard errors are clustered at the individual level. Panel A shows estimates for all stayers, comparing Network Health (solid lines) to other plans (dashed lines). Panel B shows estimates separately for stayers who are Partners patients (individuals with an outpatient visit to a Partners provider during 2011, in green) vs. all other enrollees (in blue), with solid lines continuing to denote Network Health and dashed lines other plans.

Figure 6 plots results from the event study version of (10), which also shows the empirical variation identifying the estimates. Panel A shows the overall estimates for Network Health vs. other plans (no Z_i heterogeneity). To visualize levels along with changes, I report the predicted means for Network Health ($= \exp(\overline{\alpha}_{NH} + \beta_t + \gamma_t)$) and for other plans ($= \exp(\overline{\alpha}_{Oth} + \beta_t)$), where the $\overline{\alpha}_g$'s are the constants that match the group mean in the data at the end of 2011. Costs fall sharply for Network Health stayers at the start of 2012, with a DD estimate of a 12.4% reduction (s.e. = 1.6%), or about $45 per month. By contrast, costs for other plans change very little and move in parallel to Network Health's costs aside from the one-time fall at the start of 2012. Appendix Figure 20 plots the estimates of γ_t directly, confirming the visual evidence of parallel trends (both pre and post) and suggesting that the DD estimate captures a valid causal effect.

Selection on moral hazard requires that causal reductions be larger for the types of individuals most likely to select a Partners-covering plan. Panel B of Figure 6 tests this by examining cost estimates separately by Partners patients vs. all other enrollees, the strongest predictor of selection. The graph shows two facts. First, Partners patients are much higher-cost in the pre-period (both in Network Health and other plans), consistent with them being a high-cost group. Second, Partners patients in Network Health experience much more dramatic cost reductions at the start of 2012. The DD estimate for Partners patients is a 30.6% reduction (s.e. = 3.5%), versus a reduction of just 9.2% (s.e.

= 1.7%) for other enrollees – a three-fold difference. In levels, the difference is even larger: about \$175 versus \$30 per month. Appendix Figure 21 plots the γ_t estimates, confirming the presence of parallel pre-trends and a sharp fall in 2012. After the network narrowing, Partners patient stayers in Network Health are still more costly than other stayers, but the gap has shrunk substantially: from +117% in 2011 (\$619 vs. \$285 per month) down to +40% in 2012 (\$406 vs. \$290).[34]

These results are strongly consistent with selection on moral hazard. As the theory suggests, this is natural: if use of star hospitals is concentrated among a subset of enrollees, the cost impact of dropping them should be concentrated among the same group. Appendix Table 12 (column (3)) confirms this finding in a richer specification of (10) that allow for richer Z_i heterogeneity on prior use, distance, medical risk factors, and demographics.

Appendix D.4 shows how this approach can be used to further decompose the causal effects into changes in quantity vs. price of care, following the decomposition in Section 5.1. Interestingly, about three-quarters of the causal cost reductions – including the larger reductions for Partners patients – appear to come through lower quantity, with only one-fourth coming through lower prices of care. This may reflect the importance of outpatient care (which accounts for about 70% of costs), where Partners patients appear not to use higher-price care but to utilize more intensive services (see Table 3).[35] These estimates provide further evidence that the moral hazard effect of switching to new providers involves both differences in prices for the same services and in the intensity/quantity of care provided by different doctors.

6 Structural Model

The evidence thus far suggests that plans who cover the star hospitals differentially attract high-cost consumers who are both sicker and whose costs increase more when they can access the star hospitals. In technical terms, there is adverse selection on both cost levels and on moral hazard. Moreover, the cost increases come through a mix of higher prices and higher quantities, while demand for star hospital coverage reflects both loyalty (state dependence) and persistent factors like medical needs.

What does this mix of factors imply for the incentives to cover star hospitals, and what are the associated welfare implications? This section builds on the theoretical discussion in Section 2 to estimate a structural model of demand and costs of star hospital coverage for the CommCare exchange. The model consists of two pieces: (1) a plan choice model, which draws on estimated hospital preferences as an input, and (2) an insurer cost model, which draws on the 2012 network

[34] A potential concern with this analysis is that segmenting by Partners patient status selects a temporarily sick group whose costs fall in 2012 due to mean reversion. Two findings suggest mean reversion is not driving the results. First, the use of a control group of Partners patient stayers in other plans alleviates this concern, as the DD estimate nets out any mean reversion in the control group (which does not appear to be large based on the patterns in Figure 6B). Second, a qualitatively similar pattern is apparent if I analyze enrollees by distance to Partners, which should not be subject to this concern. Costs for enrollees within 5 miles of Partners fall by 17.6% (s.e. = 3.1%), compared to a smaller fall for further enrollees of 11.1% (s.e. = 1.8%); see Appendix Figure 22 for event study estimates.

[35] Alternatively, it could reflect care disruption as patients of the dropped hospitals need to seek out new providers. The event study estimates in Appendix Figure 20 do not show much evidence that cost reductions diminish over time. But Figure 21 shows evidence that Partners patients' cost reductions may be smaller in the latter half of 2012 – about 30% versus the 40% reductions in the first half of 2012.

change for identification. Sections 6.1-6.2 present these models and their estimates. Sections 6.3-6.4 analyze the estimates to draw out the connections between cost, selection, and welfare effects of the narrower network. I use the model to analyze the incentives involved with covering the star Partners hospitals and discuss the welfare tradeoffs involved with policies that could encourage broader networks.

6.1 Insurance Plan Choice Model

I use the enrollment dataset to estimate a multinomial logit plan choice model. I treat individuals' timing of participation in the exchange as exogenous and model just their choices among exchange plans. The key assumption for my purposes is that plan changes in networks lead consumers to switch plans but do not affect whether they participate in the exchange.[36]

There are two times when plan choices are made: (1) new enrollments in the exchange (including re-enrollments after a break) and (2) plan switching decisions at annual open enrollment. I start with new/re-enrollees, who are known to make active choices, and then explain how the specification for current enrollees accounts for inertia. For new/re-enrollee i choosing at time t, the utility of plan j is:

$$U_{i,j,t}^{Plan} = \underbrace{\alpha\left(Z_i\right) \cdot Prem_{i,j,t}}_{\text{Enrollee Premium}} + \underbrace{V\left(N_{j,t}; Z_i, \beta\right)}_{\text{Network Value}} + \underbrace{\xi_{j,Reg_i,Inc_i} + \xi_{j,Reg_i,Yr_t} + \xi_j\left(Z_i\right)}_{\text{Plan Dummies (unobs. quality)}} + \underbrace{\epsilon_{i,j,t}^{Plan}}_{\text{Type 1 EV Error}} \tag{11}$$

Plan utility depends on three sets of plan attributes: premiums, networks, and unobserved quality. Subsidized enrollee premiums are observed and included directly. Networks are also observed but more difficult to capture because of their high dimensionality. To model their role, I include two sets of terms in $V(.)$. First, I follow the literature (starting with Capps et al., 2003) by including an expected "network utility" measure from a hospital choice model. Using the CommCare hospitalization sample, I estimate a multinomial logit hospital choice model and calculate the expected utility of a given plan's network $N_{j,t}$ for individual i, following the standard inclusive value formula. Appendix D.1 details the model specification, estimates, and construction of network utility. Second, I include a variable for whether the plan covers the hospitals with which the consumer has past outpatient relationships (or the share covered if there are multiple). In interpret this variable as picking up the utility of access to a hospital's physicians for outpatient care, though it may also pick up misspecification in the calculation of network utility.

The final covariates in (11) are plan j dummy variables (ξ). These both capture unobserved plan quality – e.g., insurer reputation (see Starc, 2014) – and help with the proper identification of the premium coefficient, as discussed below. There are separate plan dummies at the region x income group level (ξ_{j,Reg_i,Inc_i}) and region x year level (ξ_{j,Reg_i,Yr_t}), and I also interact plan dummies with enrollee age-sex groups and risk score deciles to allow unobserved quality to vary with medical risk. In

[36] This assumption seems reasonable because eligibility is determined by exogenous factors (e.g., income and job status) and generous subsidies encourage participation by the eligible. Further, the main variable likely to affect exchange participation, the premium of the cheapest plan, is set directly by the exchange's subsidy rules. To test the reasonableness of this assumption, Appendix Figure 19 examines whether Network Health's consumers leave the exchange at a higher rate after it narrows its network in 2012. I find no evidence of this, either overall or differentially for Partners patients or people who live near Partners.

total, the specification includes up to 77 dummy variable interactions per plan, allowing for substantial heterogeneity in unobserved quality. I can include such detailed plan dummies because premiums for the same plan vary across enrollees due to the exchange's subsidies.

A key goal of the model is to capture *heterogeneity* across consumers in price sensitivity and valuation of networks – especially heterogeneity that correlates with costs and drives adverse selection. In addition to the rich plan dummies, the model captures this by allowing coefficients on premium and network to vary with a rich set of consumer observables. I allow both premium coefficients ($\alpha(Z_i)$) and coefficients on network utility and plan coverage (β) to vary with income group, age-sex groups, immigrant status, and deciles of enrollee risk score, with the top decile split in two to allow additional flexibility for the sick. I use the HCC risk score, which is based on current-year observed diagnoses so picks up an element of unobserved risk not captured by the exchange's retrospective risk score (the "CommCare risk score").

Current Enrollees and Inertia The model so far applies to new/re-enrollees who by definition make active plan choices (otherwise, they do not get coverage). By contrast, current enrollees who fail to actively switch plans remain in their current plan by default, and this inertia is known to play a major role (Handel, 2013; Ericson, 2014). To capture inertia in a simple way, I add to the specification in (11) an extra utility for their current plan, $\delta_i \cdot 1\{j = CurrPlan_{i,t}\}$, with the coefficient δ_i allowed to vary with the same observables as the premium coefficient (income, age-sex, risk score deciles, and immigrant status). This setup ensures that the model will match average switching rates, but the coefficients themselves may pick up both true inertia and persistent unobserved heterogeneity. For my purposes, it is not clear that is important to distinguish these factors. Doing so would matter primarily for dynamic price competition, which I do not model. For robustness, I also report estimates from a specification that includes only new/re-enrollees for whom inertia is not relevant.

Identification The classic identification concern is bias from correlation between prices and unobserved quality due to strategic firm pricing. Rather than the standard practice of using instruments, I follow the alternate approach (see e.g., Nevo, 2000) of including detailed plan dummies to capture unobserved quality so that remaining price variation is plausibly exogenous. The key institutional feature is that premiums *vary across consumers for the same plan* because of CommCare's subsidies. Notably, subsidies make all plans free for enrollees with incomes below poverty, while above-poverty enrollees pay positive premiums that differ across plans. This structure makes demand patterns among below-poverty enrollees – who are unaffected by premiums – a natural "control group" for picking up shifts in unobserved plan quality.[37] Premium coefficients are estimated from how premium changes *differentially* affect above-poverty income groups' demand relative to below-poverty demand (the control group) in the same region.

Consider an example of Network Health in the Boston region from 2010-2011, as shown in Appendix Figure 10. In 2010 Network Health is the cheapest plan for all income groups, but in 2011 its premium increases (with CeltiCare replacing it as cheapest). Below-poverty groups, however, face no premium

[37] Starting in 2012 below-poverty new enrollees are limited to the choosing one of the two lowest-price plans. I account for this limitation in defining plan choice sets for these enrollees.

increase – all plans are still free – while higher income groups face differential premium increases that follow a progressive subsidy schedule. Identification of each income group's premium coefficients comes from how their demand for Network Health changes relative to the below-poverty control group. Appendix B.2 gives further information and statistics about this premium variation.

This identification strategy is analogous to difference-in-differences (DD) in a non-linear model. As in standard DD, the specification includes fixed effects to absorb all demand variation driven by endogenous price variation, leaving only the exogenous variation to identify premium coefficients. Recall that because of community rating, premiums differ only at the plan x region x year x income group level. Utility equation (11) includes plan-region-year dummies (ξ_{j,Reg_i,Yr_t}) to absorb premium differences arising from strategic pricing (since insurers set plan prices only at the region-year level). It also includes plan-region-income group dummies (ξ_{j,Reg_i,Inc_i}) to absorb persistent demand differences across income groups. All remaining variation reflects *differential premium changes* across income groups for a given plan in a given region.

The identifying assumption is that these fixed effects capture endogenous demand variation so that there are no remaining income group-specific demand trends/shocks (for a given plan in a given region) – i.e., no ξ_{j,Reg_i,Inc_i,Yr_i} – that are correlated with premium changes. One possible violation could be misspecifications in network value or inertia that (by chance) correlate with premiums. This concern suggests including additional fixed effects – specifically, the plan dummies by risk score and age-sex groups ($\xi_j\left(Z_i\right)$) that pick up any persistent differences for these groups.

A second possible violation could include advertising shocks that (by design or chance) target certain income groups. A way of testing for this is to examine whether demand trends are parallel between treatment (above-poverty) and control (below-poverty) groups around price changes. Figure 7 shows such a test. It plots monthly choice shares for new enrollees around price changes (at time 0), separately for plans that cut prices (panel A) and increase prices (panel B). Consistent with the key assumption, market shares are flat and parallel for both groups at all times except time 0. At this time, demand shifts in the expected direction for premium-paying enrollees but is unchanged for below-poverty enrollees. Thus, this analysis suggests that the fixed effects strategy is likely to generate unbiased estimates of premium coefficients.

Demand Estimates All variables entering the plan choice model are observed, so I estimate it by maximum likelihood. Appendix Tables 10-11 show the estimates. Focusing on the main summary coefficients reported in Appendix Table 10, column (2) reports the main specification including all enrollees, with inertia variables for current enrollees. Column (1) shows a robustness check with just new and re-enrollees, with inertia excluded because they make active choices. Coefficient estimates are quite similar, suggesting that the key estimates of price sensitivity and network value are robust to any challenges in identifying inertia vs. unobserved heterogeneity in plan preferences. I therefore use column (2) for the remainder of the analysis.

Focusing on column (2), premiums (in \$10 per month) enter negatively and significantly for all groups. Enrollees are quite price-sensitive: for premium-paying new/re-enrollees a \$10 per month

Figure 7: Premium Coefficient Identification: Market Shares around Price Changes

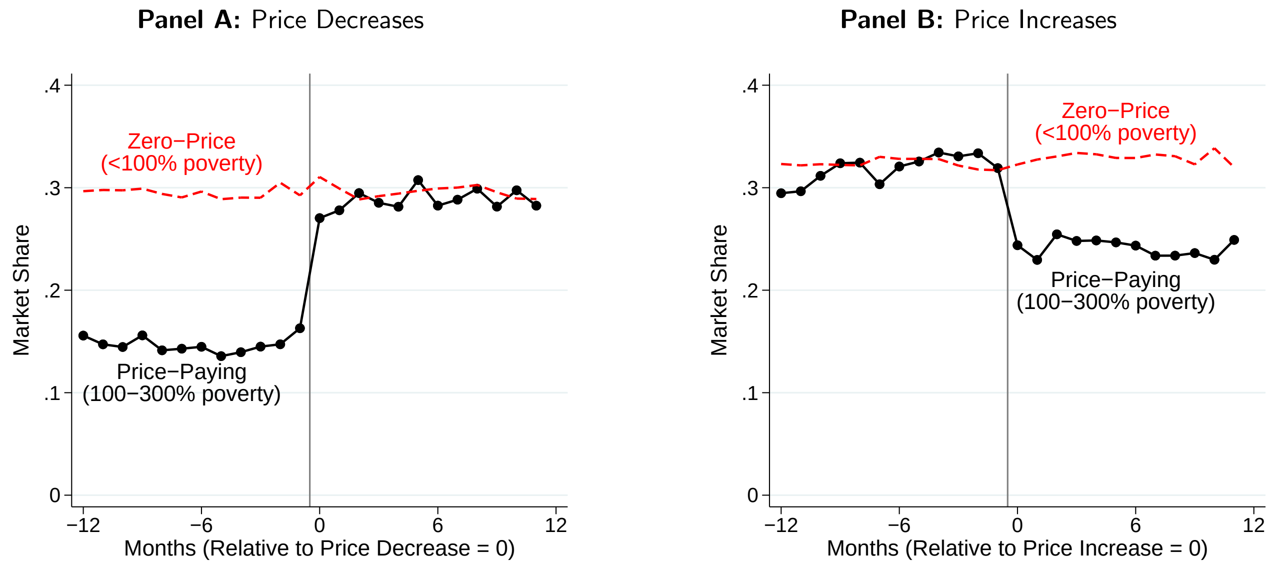

NOTE: These graphs show the source of identification for the premium coefficients in plan demand and test the key parallel trends assumption for the difference-in-differences approach. Each graph shows average monthly plan market shares among new enrollees for plans that at time 0 decreased their prices (panel A) or increased their prices (panel B). Each point represents an average market share for an independent set of new enrollees. The identification comes from comparing demand changes for above-poverty price-paying enrollees (for whom premium changes at time 0) versus below-poverty zero-price enrollees (for whom premiums are unchanged at $0). Consistent with the parallel trends assumption, trends in shares are flat and parallel for both groups at times other than the premium change but change sharply for price-payers only at the price change. The sample is limited to fiscal years 2008-2011. I drop 2012+ because below-poverty new enrollees became subject to a limited choice policy that required them to choose lower-price plans. In the demand estimates, I keep this sample but limit the choice set for this group accordingly.

premium increase lowers an average plan's market share by 26.1%. However, because enrollee premiums are low (the average is just $56.93 for above-poverty enrollees), the implied consumer-perspective demand elasticity is just -1.48, which is comparable to estimates in the literature.[38] There is substantial heterogeneity in price sensitivity, with less negative premium coefficients for higher-income, sicker, and older individuals.

As expected, there is substantial inertia in consumers' plan switching decisions, with the average coefficient of 4.413 (s.e. = 0.007). Converting inertia into dollars – by dividing each individual's inertia coefficient by their premium coefficient – implies an average "switching hurdle" of $87 per month. Though large, this estimate is actually smaller than the estimate of Handel (2013) of $2,032 per year (or $169 per month). Inertia implies that overall demand (including current enrollees) is less price elastic, with a $10 higher premium reducing market share by just 12.5% on average.

Consistent with the reduced form evidence, consumers significantly value better provider networks. This appears in both the network utility and previously used hospital variables. Network utility is normalized so that 1.0 equals the utility loss for an average Boston-area enrollee from Network Health's

[38] This is comparable to findings in the literature (see Ho (2006) for a discussion). Because of subsidies, however, the firm-perspective elasticity is much larger. A $10 price increase is a 2.5% increase relative to the average plan price of about $400. The typical firm-perspective elasticity is therefore about -10.4 (= -26.1% share change / 2.5% price change).

2012 exclusions. Narrowing the network by this magnitude reduces plan utility by an average of 0.463 (s.e. = 0.005), or about \$9.15 per month at the average premium coefficient. For people with existing provider relationships, plan utility is further reduced by 0.291 (s.e. = 0.012) on average if a plan drops all of their previously used hospitals, or \$5.75 per month at average price sensitivity. Also notable is the additional value placed by patients on coverage of *Partners* hospitals of 0.982 (s.e. = 0.021), or \$19.43 per month. As in the reduced form evidence (e.g., Figure 3B), this coefficient is consistent with consumers placing a special value on star providers.

Moreover, the estimates show substantial heterogeneity in network valuation via the interaction terms. Overall, older, sicker, and higher-income enrollees have higher utility of networks that cover their desired providers. In combination with these groups having smaller price coefficients, this implies higher willingness to pay for good provider coverage. I analyze this heterogeneity and how it relates to costs further in Section 6.4 below.

6.2 Insurer Cost Model

The second piece of the structural model is costs. The main goal of the model is to capture how expected insurer costs vary across consumers (especially based on demand for the star hospitals) and with the network change implemented by Network Health in 2012. In terms of the model in Section 2, the goal is to estimate $E\left(C_{ijt}\left(0\right)|i \in G\right)$ and $E\left(C_{ijt}\left(1\right)|i \in G\right)$ for various groups of consumers G (e.g., people with high demand for the star hospitals). Note that for the analysis below, I will restrict attention to estimating costs in a single plan (j = Network Health) in 2011-12 as it narrows its network. This avoids the need to estimate cross-plan moral hazard, which would be necessary for a full model of insurer competition.[39]

I lay out the method in two steps: (1) estimating expected costs under the plan's *observed* network ($n = 1$ in 2011 and $n = 0$ in 2012), and (2) estimating the change in costs when the network changes. Start with the former. Note that in the data we observe a consumer's *realized* costs in 2011 or 2012 under one of these networks. For instance, in $t = 2011$ we observe realized costs under the broader network (call this $C^{obs}_{ijt}\left(1\right)$). Assume that realized costs equal expected costs ($C_{ijt}\left(1\right)$) plus an idiosyncratic shock: $C^{obs}_{ijt}\left(1\right) = C_{ijt}\left(1\right) + \epsilon_{ijt}$. If the variables defining group G are known at the time when expected costs are defined (so orthogonal to the shock),[40] then $E\left(\epsilon_{ijt}|i \in G\right) = 0$ and expected costs for group G can be estimated as the average of realized costs: $\overline{C}_{G,t}\left(1\right) \equiv \frac{1}{N_G}\sum_{i\in G} C^{obs}_{ijt}\left(1\right) \rightarrow E\left(C_{ijt}\left(1\right)|i \in G\right)$ as N_G gets large. Thus, we can estimate expected costs under the actual network directly from means in the data. This method has the advantage of letting me capture cost variation in a flexible way, without relying on a parametric cost model.

The second step is estimating the *causal effect* on costs when the network changes, or $dC_i=$

[39] A previous version of this paper estimated such a model (Shepard, 2016), as do Jaffe and Shepard (2020) for the same CommCare market. However, the cross-plan moral hazard terms are (admittedly) one of the least well-identified parts of their analysis.

[40] This should be true if G is defined based on variables known prior to the realization of current-year costs (e.g., demographics, prior-observed diagnoses, or even past utilization of providers). However, because of limited availability of prior-years data (especially for new enrollees), the demand model includes the HCC risk score, which is defined using diagnoses observed in current-year claims. I therefore also need to assume that these diagnoses are known to the enrollees in advance (just not observed in the data) and are therefore exogenous.

$C_{ijt}(1)/C_{ijt}(0)$). To do so, I draw on the causal estimates of Section 5.3, which are identified from the effect of the network narrowing on stayers in Network Health from 2011-12, relative to a control group of stayers in other plans. The identification is based on a difference-in-differences logic, and Figure 6 shows evidence of parallel pre-trends. I use the estimates of Poisson regression (10) that allow for rich heterogeneity in Z_i by prior patient status (Partners and/or other dropped hospitals), distance to Partners, and the observables entering demand (income, risk score quantiles, diagnoses, and demographics). The implied causal effect of a broader network is $d\hat{C}_i = \exp(-\hat{\gamma}(Z_i))$ – with the negative sign because the estimates of $\gamma(.)$ come from the reverse experiment of a narrower network. As Section 5.3 also discusses, the approach allows for decomposing causal cost effects into quantities vs. prices. Appendix Table 12 shows the results, with columns (3)-(6) reporting estimates for insurer cost, quantity, and prices. See section 5.3 for further discussion of the estimates and their implications.

Given an estimate of either $\overline{C}_{G,t}(1)$ or $\overline{C}_{G,t}(0)$ from the data and $d\hat{C}_i$ from the regressions, I construct costs under the counterfactual network by multiplying or dividing each individual's observed costs by $d\hat{C}_i$ as appropriate. For instance, for 2011 where I observe $\overline{C}_{G,2011}(1)$, I estimate $\overline{C}_{G,2011}(0) = \frac{1}{N_G}\sum_{i\in G}\left(C^{obs}_{ijt}(1)/d\hat{C}_i\right)$. With an estimate of counterfactual costs, I can use the model to analyze the role of selection and moral hazard in the incentive to exclude star hospitals, which I turn to next.

6.3 Breakdown of Selection and Moral Hazard in 2012 Network Change

I can use the model to break down the role of selection vs. moral hazard incentives involved in Network Health's 2012 network narrowing, corresponding to the breakdown in equation (6) in the theory section. Table 4 shows the analysis. Columns (1)-(2) show observed average costs, price, and demand at the end of 2011 and start of 2012, along with the 2011-2012 change (column 3). As shown previously (Figure 1), average costs fall by 26% (\$105 per month), or 21% (\$79 per month) after risk adjustment. These average cost reductions substantially exceed the (also large) \$63 fall in the plan price, so profit margin per member increases. Because demand also increases – the number of people joining the plan because of its lower price exceeds those leaving due to the narrower network – total profits increase substantially by \$1-2 million.[41]

This analysis illustrates the strong overall incentive for the narrower network in the CommCare market. But how much of the average cost change reflects selection versus causal (moral hazard) reductions? Columns (4)-(5) of Table 4 use the model to calculate two counterfactuals: holding fixed 2011 enrollment but using the model to predict costs under the narrower 2012 network (column 4) and applying the broader network costs to the 2012 enrollees (column 5). Comparing columns (1)/(4) and (2)/(5) yields estimates of *moral hazard reductions* (cost change with fixed enrollment), while comparing columns (1)/(5) and (2)/(4) yields estimates of the *selection effect* (changing enrollment, holding the network fixed).

The estimates (shown in the bottom of the table) suggest both effects are sizable. Holding enroll-

[41] Two caveats are worth noting. First, this is a measure of gross profits before administrative cost, which I do not observe in the claims data. Second, these outcomes are a function of both Network Health's and its competitors actions, as well as the limited choice policy change in 2012. While other plans do not meaningfully change networks, prices do change as shown in Appendix Figure 9. Results are similar if I limit the analysis to either the above-poverty population (not subject to the limited choice policy) or the below-poverty population (who do not pay prices).

ment fixed, causal moral hazard reductions lower costs by \$48-60 per month, which is 46-57% of the fall in raw average costs and 61-75% of the fall in risk-adjusted average costs. The remainder of the fall comes from the selection effect. Thus, moral hazard represents the bulk of the cost changes, but selection plays an important role (25-39% of the change) even after risk adjustment. The table also illustrates the finding of "selection on moral hazard." The moral hazard estimate is larger for the 2011 population (who chose the plan with its broader network) than for 2012 enrollees (who select it with its narrower network). The difference is a meaningful \$11 per month, or about one-fifth larger moral hazard.

Overall, Table 4 illustrates the strong economic forces pushing towards excluding expensive star hospitals in a low-income insurance market. Even if selection incentives were completely eliminated, moral hazard cost reductions of \$48-60 (or about 15%) are enormous in a market where consumers are price elastic enough that each \$10 lower premium raises demand by 26%. Moreover, as I show in the next subsection, \$48-60 is larger than consumer WTP for star hospital coverage for all but the top 3-5% of consumers. The plan therefore has substantial ability to compensate most consumers for the narrower network with a lower price, leaving both parties better off. Moreover, selection further increases the incentive to narrow the network. Even if the star hospitals dramatically cut their prices (or the state offered subsidies to cover its higher costs) in a way that eliminated the moral hazard effect of covering them, the \$24-31 risk-adjusted selection effect is larger than WTP for all but the top 11-15% of consumers. Thus, the plan could exclude the star hospitals and cut its price, leaving most of its consumers happier while the (expensive) consumers unhappy with the change switch to another plan.

6.4 Analysis: Cost, Selection, and Welfare Impacts of Star Hospital Coverage

I next analyze WTP and costs under the broader 2011 vs. narrower 2012 networks in the style of Einav, Finkelstein, and Cullen (2010a, "EFC").[42] This approach provides a useful way of summarizing demand/cost primitives to understand the forces driving adverse selection and welfare.[43] It works by ranking consumers in terms of decreasing WTP types for the broader network (call this ranking $s \in [0,1]$, with $s = 0$ having highest WTP and $s = 1$ lowest) and plotting WTP and costs for the average consumer in each s bin. The key variable is WTP for the broader network, defined based on the plan utility estimates of equation (11):

$$\Delta WTP_i \equiv \tfrac{1}{-\alpha(Z_i)} \cdot \left[V\left(N_{NH,2011,}; Z_i, \beta\right) - V\left(N_{NH,2012}; Z_i, \beta\right)\right] \tag{12}$$

[42]Although this change involves more than just the star Partners hospitals, Partners comprises the large majority of the dropped hospital capacity and has the largest patient demand. Partners hospitals comprise 76% of the 3,207 hospital beds in the dropped hospitals. Partners patients comprise 67% of the switchers out of Network Health in 2012 (vs. 8% patients at other dropped hospitals).

[43]EFC demand/cost curves are also sufficient for equilibrium in a stylized setting where perfectly competitive (and otherwise identical) insurers compete on a single binary quality variable – here, broad vs. narrow network. Of course, real-world markets are much more complex, but the EFC curves are remarkably useful in capturing the underlying economic forces. A previous version of this paper modeled a more realistic (but still stylized) version of imperfect competition in the CommCare market and found results in line with the discussion below (Shepard, 2016).

Table 4: Analysis of Selection vs. Moral Hazard Incentives for 2012 Network Change

	Actual Enrollees/Networks			Counterfactuals	
Enrollees:	2011 Enr.	2012 Enr.	*Change*	2011 Enr.	2012 Enr.
Network:	Broader	Narrower	*2011-12*	Narrower	Broader
	(1)	(2)	(3)	(4)	(5)
Average Cost	$403.8	$299.1	*-$104.8* *[-26%]*	$344.4	$347.4
Risk-Adj. Avg. Cost	$377.5	$298.4	*-$79.2* *[-21%]*	$321.9	$346.6
Plan Price	$422.6	$360.0	*-$62.6*		
Demand (memb./month)	46,880	50,589	*+7.9%*		
Profit before risk adj. ($milllion)	$0.88	$3.08	*+$2.20*		
after risk adjustment	$2.11	$3.12	*+$1.00*		
Implied Selection and Moral Hazard:					
Moral Hazard Reductions:				*-$59.5*	*-$48.4*
Selection Effect: Raw Costs				*-$45.3*	*-$56.4*
Risk-Adj Cost				*-$23.6*	*-$30.9*

NOTE: The table breaks down Network Health's large fall in average costs in 2012 (see Figure 1) into components driven by selection and moral hazard (causal cost reductions), following the decomposition in equation (6) in the theory section. Columns (1)-(2) show Network Health's observed costs (and price and demand) in the final quarter of 2011 and first quarter of 2012, with column (3) showing the change between these two. Column (4) shows the counterfactual of costs for the 2011 enrollees under the narrower (2012) network; column (5) shows the counterfactual of costs for 2012 enrollees under the broader (2011) network. Both of these cost counterfactuals are predicted using the cost model estimated in Section (6.2).

where $V(.)$ is the consumer's network value terms for the 2011 and 2012 network. As is standard, WTP equals the change in network value divided by -1 times the premium coefficient (utility of money).[44]

The other key variables are costs. For simplicity, I focus on enrollees in Network Health in 2011; results are similar if I examine other groups such as enrollees in 2012. Costs are defined as described in Section 6.2: using observed 2011 costs for the broad network and scaling these down by the model's estimated cost change ($d\hat{C}_i$) for the narrower network. I then plot cost variables (levels and differences) conditional on WTP ranking s – e.g., $\overline{C}(1;s) = E(C_{ijt}(1)|s)$ for costs under the broad network – which correspond to type-specific (or "marginal") cost curves in the EFC framework. Consistent with the exchange definition, risk-adjusted costs equal average cost divided by average (CommCare) risk score for a given s.

Figure 8 shows results. Each point on the graphs represents a 2-percentile bin of the (decreasing) WTP distribution up to the 40th percentile; after this WTP is quite low so I summarize the remaining distribution a single point. Panel A shows the cost curves. Two results stand out. First, there is substantial adverse selection for costs under the broad network, indicated by steeply downward sloping cost curves with decreasing WTP type. Applying risk adjustment (light-blue dashed curve) makes a

[44] I do not have $\alpha(Z_i)$ estimate for below-poverty enrollees, so for them I use the α estimates for comparable 100-150% of poverty enrollees. This may overstate WTP (since α is generally more negative for poorer people) which is conservative given my findings of low WTP.

large difference, but costs are still steeply downward sloping for the broad network. Risk-adjusted costs in the top-2% WTP bin are \$628 per month, about 50% larger than at the 20th percentile (\$416) and twice the cost at the 40th percentile (\$309).

Second, and by contrast, risk-adjusted costs under the *narrower network* (red dashed curve) are much flatter. Except for the top 2% point (\$434), the curve is relatively flat in the \$280-360 range. Put differently, most of the risk-adjusted selection comes from the larger incremental costs for high-WTP types, which is reflected in the larger gap between the two dashed cost curves for high-WTP types. This again illustrates the selection on moral hazard pattern in the estimates.

Panel B of Figure 8 shows this result more directly and plots the key curves for a standard welfare analysis: $\Delta WTP\,(s)$ and incremental costs (moral hazard) $\Delta Cost\,(s) \equiv \overline{C}\,(1; s) - \overline{C}\,(0; s)$. Incremental costs are downward sloping with WTP and *everywhere above* the WTP curve by a factor of 3-6x throughout the distribution. Thus, under a standard surplus measure ($\Delta Surplus = \Delta WTP - \Delta Cost$), coverage of the broader network including the star hospitals is inefficient, as insurer costs exceed consumer valuations. This holds true throughout the WTP distribution because of the way $\Delta Cost$ rises steeply with WTP. On average, WTP for the broader network is \$11 per month versus average $\Delta Cost$ of \$58 per month. But even though people in the top 2% of WTP place substantially higher value on the network – about \$90 per month, or almost twice the average enrollee premium in CommCare[45] – their incremental costs are even larger (\$361 per month). Indeed, because $\Delta Cost$ is *steeper* than ΔWTP, the $\Delta Surplus$ measure is actually most negative for the highest-WTP types – consistent with the "backward sorting" pattern found by Marone and Sabety (2019). The people who demand Partners coverage the most are (under a standard welfare metric) the people for whom it is *least* efficient.

It is important to emphasize that policymakers may care about factors beyond standard market surplus in judging social welfare and deciding whether to subsidized/mandate coverage of star hospitals. Nonetheless, the basic finding that costs of star hospital coverage are large relative to consumer valuation appears robust. Appendix D.5 presents ΔWTP and $\Delta Cost$ under several modifications. These include counting only quantity of care reductions in $\Delta Cost$, recalculating $\Delta Cost$ using 10-50% lower Partners prices (reflecting either Partners' markups or price reductions), and redefining ΔWTP based on a lower social marginal utility of money ($-\alpha$ in (12)). Cost continues to exceed WTP in all of these modifications individually. Only a combination of (extremely large) 50% Partners price cuts and scaled up ΔWTP (using the 99th percentile lowest $-\alpha$) can close the gap between WTP and costs.

The WTP and cost curves shown here reflect preferences for using star hospitals inferred from choices in the data. As emphasized elsewhere, these preferences are partly state dependent: they are driven by patient-doctor relationships whose strength increases with recent use and may fade way over time with lack of use. How does this affect the welfare analysis? State dependence creates a divergence

[45] For further context, Finkelstein et al. (2019) find that median WTP for insurance overall relative to uninsurance is about \$100 per month, so a \$90 value for a broader network is quite large. Ericson and Starc (2015b) study a higher-income Massachusetts population and find that *typical* WTP for a broad network (that includes Partners) vs. a narrower network (that excludes Partners) is between \$68-123 per month. This is comparable to the highest-WTP types in CommCare's low-income population and much higher than the average WTP of \$11 per month or the median of \$4.7 per month.

Figure 8: Cost and Willingness to Pay Curves for Broader Network

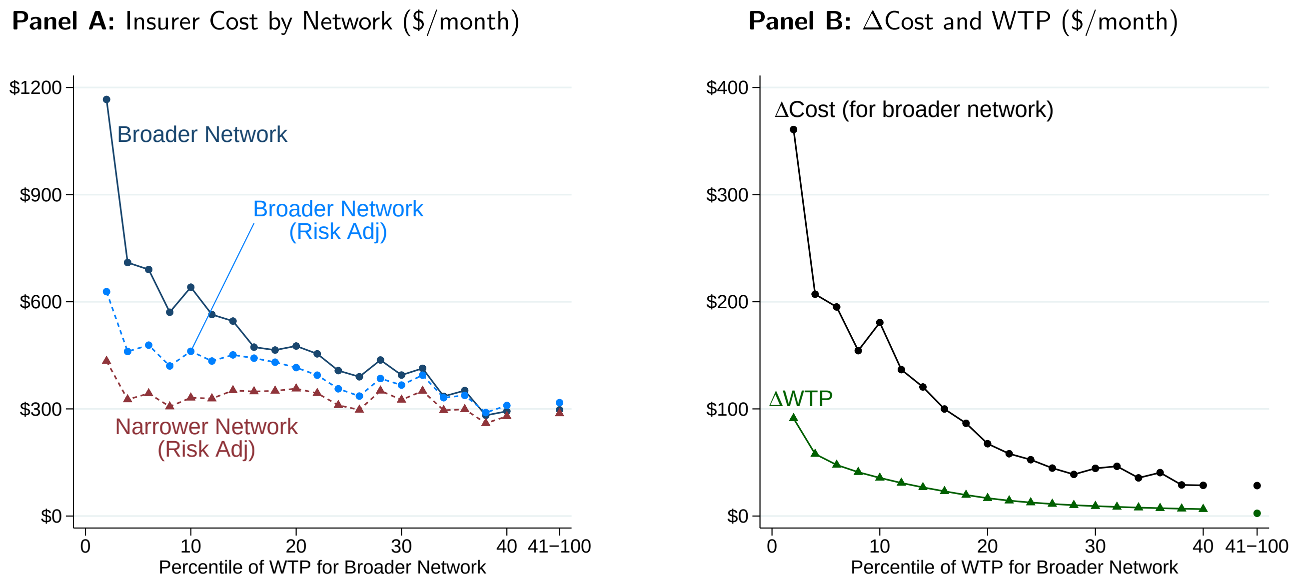

NOTE: These graphs show cost and willingness to pay (WTP) curves derived from the structural model estimates. The x-axis for both panels is the WTP type (s), the percentile ranking of WTP for Network Health's broader 2011 network that includes the star Partners hospitals, relative to the narrower 2012 network that excludes Partners. WTP declines moving left to right. Panel A shows type-specific raw insurer costs under the broader network (solid dark blue), risk-adjusted costs under the broad network (dashed light blue), and risk-adjusted costs under the narrow network (dashed red). The downward slope of these curves indicates adverse selection. Panel B shows the type-specific incremental cost (moral hazard) of the broader network ($\Delta Cost$) and the ΔWTP for the network. $\Delta Cost$ slopes down steeply (consistent with selection on moral hazard) and is everywhere above WTP (consistent with negative surplus of the broader network).

between short- and long-run implications of losing access to a preferred good. Normally, switching costs imply that short-run utility losses (as reflected in ΔWTP curve shown) are *larger than* long-run losses. This would tend to reinforce the finding that WTP falls short of costs. However, in this case, $\Delta Cost$ is *also* driven by preferences for using star hospitals. In the long-run, a patient who loses star hospital access and develops new provider relationships will also generate lower $\Delta Cost$ if they regain star hospital access. Thus, it is unclear how state dependence affects the long-run net benefit (WTP minus costs) of star hospital coverage. The change will depend on whether $\Delta Cost$ or ΔWTP fall more in the long-run.

Discussion of Policy Implications The finding that $\Delta Cost$ rises with (and exceeds) ΔWTP for Partners coverage throughout the distribution indicates the strength of selection on moral hazard and suggests the policy challenge involved. At any single premium for a Partners-covering plan, the people who select into the plan will be individuals with the largest moral hazard ($\Delta Cost$) and whose own expected moral hazard exceeds their own WTP for Partners coverage. While modifications to risk adjustment or a sufficiently large subsidy can ensure that insurers are willing to cover Partners, these standard policies cannot change this basic underlying welfare reality. Risk adjustment and subsidies are best designed to address selection on baseline cost levels (i.e., $C_{ij}(0)$) not selection on moral hazard. Moreover, these policies would run the risk of exacerbating Partners' existing bargaining power, leading to even higher prices and costs. An alternate policy would be to allow insurers to vary incremental premiums with factors like distance correlated with expected Partners use (Bundorf et al., 2012), but this would likely fall short given that $\Delta Cost > \Delta WTP$ throughout the distribution.

More promising would be policies that directly addressed the moral hazard involved in high $\Delta Cost$. The most natural policy would be allowing higher "tiered" copays when patients use expensive providers (Prager, 2018), a policy disallowed in CommCare and Medicaid markets. This works both by reducing insurers' $\Delta Cost$ directly for high-Partners use patients (since insurers pay less reflecting the tiered copay) and by leading patients to shift care to lower-cost providers. Importantly, to work well patients would need to specifically reduce *low-value* use of Partners (use for which the the value falls short of the copay); otherwise, ΔWTP would also shift down analogously. This is what standard economic theory predicts should occur, though some evidence on patient choice calls into question this prediction (e.g., Brot-Goldberg et al., 2017).

7 Conclusion

As public programs increasingly use markets for health insurance, an important question is how well insurance competition will work. A key aspect of this question is whether adverse selection is still important, despite policies intended to combat it. This paper shows evidence from Massachusetts' pioneer exchange that even with sophisticated risk adjustment, selection creates a significant disincentive to covering the state's most prestigious star hospitals. This occurs partly through a mechanism that, while intuitive, has not previously been highlighted. People select plans based on their preferences for the star hospitals. And these consumers have high costs precisely because they use the expensive star providers for care. This creates selection on a dimension of costs unlikely to be offset by medical risk adjustment.

Although these results are from a specific setting, they have general implications. The mechanism I highlight is general: there are high-price star hospitals across the country (Ho 2009) and patients surely vary in their preferences for them (e.g., based on distance and past relationships). Therefore, adverse selection is likely to emerge in markets like the ACA exchanges. My findings may help explain the sharp rise of narrow networks, which tend to exclude star hospitals. The findings also suggest that star hospitals may face a more challenging economic environment as market-based insurance expands both in public programs (via the ACA and Medicare Advantage) and employer insurance (via private exchanges). Star hospitals may face the choice of either lowering prices or losing access to a large group of patients.

The findings also have general implications for how economists think about adverse selection in health insurance markets. My results suggest that consumer preferences for high-cost treatment options – star hospitals in my study, but the same idea could apply to any expensive provider, drug, or treatment – can naturally lead to adverse selection, and specifically selection on moral hazard. Selection on moral hazard is not just an empirical curiosity but affects welfare and policy implications. Typically, economists think of adverse selection as leading to too little access to (or enrollment in) generous insurance, creating a rationale for mandates or subsidies. But selection on moral hazard complicates the analysis because people with the greatest demand for a generous benefit also have the largest cost increases from it. As a result, subsidies for generous coverage may not improve welfare, as indicated by my model estimates that WTP for star hospital coverage falls far short of its moral hazard costs for most consumers.

These results suggest the importance of distinguishing selection on cost levels vs. moral hazard in future empirical work. They also show the importance of studying alternate policies to address these inefficiencies. Fundamentally, these problems are related to a basic sorting challenge: which patients should get access to the expensive services star hospitals provide? In the current system, consumers get access to star hospitals based on their plan choice, after which use of these providers is highly subsidized by the insurer. This setup leads to higher costs (moral hazard) and selection on moral hazard. Policies that reduce this moral hazard – e.g., higher tiered copays for expensive hospitals or incentives for doctors to refer patients more efficiently – may also mitigate the adverse selection. However, these policies need to be balanced against potential losses to risk protection and access to star hospitals. Better understanding the optimal balance is an important topic for future work.

References

Eduardo M Azevedo and Daniel Gottlieb. Perfect competition in markets with adverse selection. *Econometrica*, 85(1):67–105, 2017.

Friedrich Breyer, M Kate Bundorf, and Mark V Pauly. Health care spending risk, health insurance, and payment to health plans. In *Handbook of health economics*, volume 2, pages 691–762. Elsevier, 2011.

Zarek C Brot-Goldberg, Amitabh Chandra, Benjamin R Handel, and Jonathan T Kolstad. What does a deductible do? the impact of cost-sharing on health care prices, quantities, and spending dynamics. *The Quarterly Journal of Economics*, 132(3):1261–1318, 2017.

Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does risk selection respond to risk adjustment? new evidence from the medicare advantage program. *American Economic Review*, 104(10):3335–64, 2014.

M Kate Bundorf, Jonathan Levin, and Neale Mahoney. Pricing and welfare in health plan choice. *The American Economic Review*, 102(7):3214–3248, 2012.

Cory Capps, David Dranove, and Mark Satterthwaite. Competition and market power in option demand markets. *RAND Journal of Economics*, pages 737–763, 2003.

Colleen Carey. Technological change and risk adjustment: Benefit design incentives in medicare part d. *American Economic Journal: Economic Policy*, 9(1):38–73, 2017.

David Chan and Jonathan Gruber. How sensitive are low income families to health plan prices? *The American Economic Review*, 100(2):292–296, 2010.

Amitabh Chandra, Jonathan Gruber, and Robin McKnight. The impact of patient cost-sharing on low-income populations: evidence from massachusetts. *Journal of health economics*, 33:57–66, 2014.

(CHIA) Massachusetts Center for Health Information and Analysis. Annual report on the performance of the massachusetts health care system. Technical report, https://www.chiamass.gov/annual-report/, 2014a.

(CHIA) Massachusetts Center for Health Information and Analysis. Massachusetts hospital profiles: Data through fiscal year 2012. Technical report, https://www.chiamass.gov/hospital-profiles/, 2014b.

Pierre-André Chiappori and Bernard Salanie. Testing for asymmetric information in insurance markets. *Journal of political Economy*, 108(1):56–78, 2000.

Gary Claxton, Matthew Rae, Anthony Damico, Gregory Young, Daniel McDermott, and Heidi Whitmore. Health benefits in 2019: premiums inch higher, employers respond to federal policy. *Health Affairs*, 38(10):1752–1761, 2019.

Jeffrey Clemens and Joshua D Gottlieb. In the shadow of a giant: Medicare's influence on private physician payments. *Journal of Political Economy*, 125(1):1–39, 2017.

Martha Coakley. Massachusetts attorney general investigation of health care cost trends and cost drivers: Preliminary report. Technical report, Office of Mass. Attorney General, https://www.mass.gov/doc/2010-investigation-of-health-care-cost-trends-and-cost-drivers-preliminary-report/download, 2010.

Zack Cooper, Stuart V Craig, Martin Gaynor, and John Van Reenen. The price ain't right? hospital prices and health spending on the privately insured. *The Quarterly Journal of Economics*, 134(1): 51–107, 2019.

Vilsa Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya. Can health insurance competition work? evidence from medicare advantage. Technical report, National Bureau of Economic Research, 2014.

David M Cutler and Sarah J Reber. Paying for health insurance: the trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466, 1998.

Leemore S Dafny, Igal Hendel, Victoria Marone, and Christopher Ody. Narrow networks on the health insurance marketplaces: prevalence, pricing, and the cost of network breadth. *Health Affairs*, 36(9): 1606–1614, 2017.

Francesco Decarolis and Andrea Guglielmo. Insurers response to selection risk: Evidence from medicare enrollment reforms. *Journal of health economics*, 56:383–396, 2017.

Sandra L Decker. Two-thirds of primary care physicians accepted new medicaid patients in 2011–12: a baseline to measure future acceptance rates. *Health affairs*, 32(7):1183–1187, 2013.

Liran Einav, Amy Finkelstein, and Mark R Cullen. Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics*, 125(3):877–921, 2010a.

Liran Einav, Amy Finkelstein, and Jonathan Levin. Beyond testing: Empirical models of insurance markets. *Annual Review of Economics*, 2(1):311–336, 2010b.

Liran Einav, Mark Jenkins, and Jonathan Levin. Contract pricing in consumer credit markets. *Econometrica*, 80(4):1387–1432, 2012.

Liran Einav, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, 2013.

Liran Einav, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf. Beyond statistics: the economic content of risk scores. *American Economic Journal: Applied Economics*, 8(2):195–224, 2016.

Randall P Ellis and Thomas G McGuire. Predictability and predictiveness in health care spending. *Journal of health economics*, 26(1):25–48, 2007.

Keith M Marzilli Ericson. Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64, 2014.

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

Keith M Marzilli Ericson and Amanda Starc. How product standardization affects choice: Evidence from the massachusetts health insurance exchange. *Journal of Health Economics*, 50:71–85, 2016.

Keith Marzilli Ericson and Amanda Starc. Measuring consumer valuation of limited provider networks. *American Economic Review*, 105(5):115–19, 2015b.

Yevgeniy Feyman, Jose F Figueroa, Daniel E Polsky, Michael Adelberg, and Austin Frakt. Primary care physician networks in medicare advantage. *Health Affairs*, 38(4):537–544, 2019.

Amy Finkelstein, Nathaniel Hendren, and Mark Shepard. Subsidizing health insurance for low-income adults: Evidence from massachusetts. *American Economic Review*, 109(4):1530–67, 2019.

Richard G Frank, Jacob Glazer, and Thomas G McGuire. Measuring adverse selection in managed health care. *Journal of Health Economics*, 19(6):829–854, 2000.

M Gaynor and WB Vogt. Competition among hospitals. *The Rand journal of economics*, 34(4): 764–785, 2003.

Michael Geruso and Timothy J Layton. Selection in health insurance markets and its policy remedies. *Journal of Economic Perspectives*, 31(4):23–50, 2017.

Michael Geruso and Thomas G McGuire. Tradeoffs in the design of health plan payment systems: Fit, power and balance. *Journal of Health Economics*, 47:1–19, 2016.

Michael Geruso, Timothy Layton, and Daniel Prinz. Screening in contract design: evidence from the aca health insurance exchanges. *American Economic Journal: Economic Policy*, 11(2):64–107, 2019.

Soheil Ghili. Network formation and bargaining in vertical markets: The case of narrow networks in health insurance. *Available at SSRN 2857305*, 2016.

Sherry Glied. Managed care. In *Handbook of health economics*, volume 1, pages 707–753. Elsevier, 2000.

Gautam Gowrisankaran, Aviv Nevo, and Robert Town. Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203, 2015.

Jonathan Gruber and Robin McKnight. Controlling health care costs through limited network insurance plans: Evidence from massachusetts state employees. *American Economic Journal: Economic Policy*, 8(2):219–50, 2016.

Benjamin R Handel. Adverse selection and inertia in health insurance markets: When nudging hurts. *The American Economic Review*, 103(7):2643–2682, 2013.

Health Care Cost Institute. 2014 health care cost and utilization report: Analytic methodology. 2015.

Lucas Higuera, Caroline S Carlin, and Bryan Dowd. Narrow provider networks and willingness to pay for continuity of care and network breadth. *Journal of health economics*, 60:90–97, 2018.

Kate Ho and Robin S Lee. Insurer competition in health care markets. *Econometrica*, 85(2):379–417, 2017.

Kate Ho and Robin S Lee. Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, 109(2):473–522, 2019.

Kate Ho and Ariel Pakes. Hospital choices, hospital prices, and financial incentives to physicians. *American Economic Review*, 104(12):3841–84, 2014.

Katherine Ho. The welfare effects of restricted hospital choice in the us medical care market. *Journal of Applied Econometrics*, 21(7):1039–1079, 2006.

Katherine Ho. Insurer-provider networks in the medical care market. *American Economic Review*, 99 (1):393–430, 2009.

Gretchen Jacobson, Ariel Trilling, Tricia Neuman, Anthony Damico, and Marsha Gold. Medicare advantage hospital networks: how much do they vary. *Kaiser Family Foundation*, 2016.

Sonia Jaffe and Mark Shepard. Price-linked subsidies and imperfect competition in health insurance. *American Economic Journal: Economic Policy*, 12(3):279–311, 2020.

Kurt Lavetti and Kosali Simon. Strategic formulary design in medicare part d plans. *American Economic Journal: Economic Policy*, 10(3):154–92, 2018.

Eli Liebman. Bargaining in markets with exclusion: An analysis of health insurance networks, 2016.

Victoria R Marone and Adrienne Sabety. Should there be vertical choice in health insurance markets? Technical report, Working Paper, 2019.

Center for U.S. Health System Reform McKinsey. Hospital networks: Perspective from four years of the individual market exchanges. Technical report, https://healthcare.mckinsey.com/sites/default/files/20172017.

Aviv Nevo. Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, pages 395–421, 2000.

Joseph P Newhouse. Accounting for teaching hospitals' higher costs and what to do about them. *Health Affairs*, 22(6):126–129, 2003.

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

Elena Prager. Consumer responsiveness to simple health care prices: Evidence from tiered hospital networks, 2018.

Devesh Raval and Ted Rosenbaum. Why do previous choices matter for hospital demand? decomposing switching costs from unobserved preferences. *Review of Economics and Statistics*, 100(5):906–915, 2018.

Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.

Adrienne Sabety. The value of relationship-specific capital in health care. Technical report, Working Paper, 2020.

Mark Shepard. Hospital network competition and adverse selection: evidence from the massachusetts health insurance exchange. Technical report, NBER Working Paper 22600, 2016.

Mark Shepard and Myles Wagner. Automatic enrollment in health insurance: Evidence and policy tradeoffs. Technical report, Working paper, 2020.

Zirui Song, Dana Gelb Safran, Bruce E Landon, Mary Beth Landrum, Yulei He, Robert E Mechanic, Matthew P Day, and Michael E Chernew. The 'alternative quality contract,' based on a global budget, lowered medical spending and improved quality. *Health Affairs*, 31(8):1885–1894, 2012.

Amanda Starc. Insurer pricing and consumer welfare: Evidence from Medigap. *The RAND Journal of Economics*, 45(1):198–220, 2014.

Robert Town and Gregory Vistnes. Hospital competition in hmo networks. *Journal of health economics*, 20(5):733–753, 2001.

André Veiga and E Glen Weyl. Product design in selection markets. *The Quarterly Journal of Economics*, 131(2):1007–1056, 2016.

Jacob Wallace, Anthony Lollo, and Chima D Ndumele. Comparison of office-based physician participation in medicaid managed care and health insurance exchange plans in the same us geographic markets. *JAMA network open*, 3(4):e202727–e202727, 2020.

Online Appendix:

Hospital Network Competition and Adverse Selection

Mark Shepard

Appendix A: Data Construction and Summary Statistics

A.1 Hospitalization Dataset Construction

To estimate the hospital choice and prices model, I use the CommCare insurance claims data to construct a dataset of enrollees' inpatient hospitalizations at acute care hospitals in Massachusetts. Constructing hospital visits from claims data involves extensive cleaning. I base my procedure on a method used by the Health Care Cost Institute (Health Care Cost Institute, 2015; see also Cooper et al., 2019), modified to my setting and the nature of the CommCare insurer claims.

I start by flagging inpatient hospital facility claims, based on having a valid site of service code[46] plus either a valid revenue code for "room and board" services[47] or a valid DRG code. I further restrict to claims where the billing provider is a Massachusetts acute care hospital, which excludes out-of-state hospitals (relatively rare, but for which I do not have network information) and inpatient stays at skilled nursing facilities, psychiatric hospitals, and rehab hospitals (many of which are also for mental health/substance abuse, which is quite common in the CommCare data).[48] I do retain claims for several prominent specialty hospitals: New England Baptist (orthopedics), Mass Eye & Ear Infirmary, Dana Farber Cancer Institute, and Boston Children's Hospital. However, these are relatively uncommon (<1% of admissions combined).

Using this dataset of inpatient hospital facility claims, I define inpatient "episodes," which includes all consecutive days when a patient is hospitalized. This sometimes includes multiple adjacent admissions (typically when a patient is transferred), which I will subsequently split out. I group together all adjacent/overlapping inpatient hospital facility claims based on the admission and discharge dates on the claims.[49] Using this episode sample, I then add on *all* claims (including professional and ancillary services) that occurred on a day the patient was admitted.[50] I also include emergency department

[46]The inpatient site of service codes are: (for the UB-04 bill type) U11, U12, U15, U16, U18, and (for CMS-1500 bill type) C21.

[47]Specifically, these include: all-inclusive codes 100-101; room and board codes 110-159, excluding the codes for hospice and rehabilitation; and ICU and CCU codes 200-219. I do not include newborn nursery codes, since all CommCare enrollees are adults.

[48]I define providers using a hand-constructed dataset made from the provider name, type, and location reported on the claims' provider file.

[49]These dates typically make sense and are consistent within claims for a hospitalization. But some hospitals appear to submit multiple adjacent-dated claims for each hospitalization (e.g., one claim per day, with admit date = discharge date). This procedure groups these together into a single admission. As a safeguard, I drop a tiny number of episodes (0.01%) where this extends the implied length of stay by more than 14 days.

[50]I exclude a small number of claim lines (0.3%) added via this procedure that occur at non-acute hospitals. These are often claims for a post-acute/rehab stay that begins the day of discharge.

(ED) and ED observation visits that occur the day prior to admission.[51]

From this dataset of all claims for a hospitalization episode, I collapse the data to the hospitalization level. I calculate insurer payment and patient cost sharing amounts by summing across all claim lines – both total and separately for inpatient facility claims, professional services, and outpatient facility claims (typically ED visits). I define the principal diagnosis using the primary (first) diagnosis code associated with the main inpatient facility claims for the hospitalization.[52] For my model I categorize principal diagnoses into Clinical Classifications Software (CCS) codes – a useful grouping defined by the U.S. Agency for Healthcare Research and Quality (AHRQ) that collapses detailed ICD-9 codes into about 280 clinically meaningful categories. I define comorbidities using dummy variables for Elixhauser categories – based on whether an associated diagnosis code appears as a primary or secondary diagnosis on any of the claim lines for the hospitalization. I define the DRG using the value reported on the inpatient facility claims when available (86% of episodes).[53] These reported DRGs are mostly MS-DRGs version 25, though versions 23-24 and APR-DRGs also appear on the data. Since my goal is to have a consistent service unit measure for inpatient pricing (see Appendix C.2), I either map earlier-version MS-DRGs to version 25 (where the match is appropriate) or into a unique DRG category (to avoid a false overlap with version 25).[54] In the 14% of cases with no reported DRG, I leave the DRG as missing and instead use the CCS code of the principal diagnosis as the service unit for the hospital price model.

Finally, I limit the sample in several ways to facilitate estimation and exclude admissions where the data may be incorrect. Starting from a sample of 81,179 episodes, I exclude 1,780 (2.2% of the sample) where the episode included admissions at multiple different hospitals; in these cases (which are likely transfers), the patient choice is ambiguous. I further exclude 1,245 episodes (1.5% of the sample) where the total facility paid amount is <\$100 (most of these are \$0); these are likely either errors, denied claims, or corner cases where my data cleaning procedure fails to work properly. Next, I exclude 2,184 admissions from FY 2007 (for which I do not have network information), 5,552 episodes from FY 2014 (which is outside my sample period of interest), and 2 admissions that lack both DRG and principal diagnosis information. Finally, I exclude admissions where the patient zip code is missing/invalid (17 cases, 0.02% of the sample) or the patient used a hospital more than 100 miles away (305 cases, 0.39% of the sample). The latter is a standard restriction in empirical hospital choice models that lets me keep the choice set size manageable. The final hospitalization dataset includes 70,094 hospitalizations

[51] Following HCCI, ED claims are identified by including a line with associated revenue codes (450-452, 456, 459, or 981) or procedure (HCPCS) codes for E&M services in the ED (99281-99292, 99466-99476). Observation stays are identified by revenue codes (760-762, or 769) or HCPCS procedure codes (99217-99220). I also use the ED claim line definition to flag whether a hospitalization was for an emergency, based on including an ED visit.

[52] The vast majority (~90%) of hospitalizations have a single inpatient facility claim. In the remaining cases where there are multiple claims, I use the diagnosis associated with the highest total paid amounts on facility claims for the episode.

[53] In about 2% of cases, there are multiple reported DRGs. In these cases, I use the DRG associated with the inpatient claim with the highest total paid amounts.

[54] To do the mapping, I use the DRG code listed on claims when either: (1) the hospital-insurer pair pays using version 25, or (2) the hospital-insurer pair uses v23 or v24 and the DRG code definition is consistent between these versions and v25. In remaining cases, I map the DRG on the claims as a unique code, making sure it does not accidentally map to an existing v25 code. After doing this procedure, most admissions (about 74%) map to MS-DRG v25. Another 24% are version 24, and there are also a few from v23 (~1%), APR-DRG (~1%), and unknown values (0.3%).

over the FY 2008-2013 period.

A.2 Outpatient Care Provider Use Dataset Construction

As described in Section 3.2, I construct a dataset of whether enrollees have used certain hospitals or their affiliated community health centers (CHC) for outpatient care. Starting from the full claims data, I exclude inpatient and emergency department care, following a similar definition as in the hospitalization dataset. Emergency department care is defined in the same way as for the hospitalization file (see Appendix A.1 above). Inpatient care is flagged based on having either a valid inpatient site of service code, a valid revenue code for "room and board" services or a valid DRG code. This definition is slightly broader than for the hospitalization dataset in that it counts care as inpatient based on the site of service code alone. My goal is to be conservative and avoid including inpatient care in my outpatient care file. After excluding these inpatient/ED claims, I limit to outpatient and professional services using a flag given by the data provider.

I code the hospital or CHC (if any) at which the outpatient care was delivered using the name of the billing provider on the claims. This process involved hand-cleaning the names on the insurance provider file. By using the billing provider, I capture services delivered by physicians employed by a hospital or treating at a hospital-owned practice. This is intentional, since these physicians are closely associated with the hospital and are excluded from network in the change I study. I link CHCs to hospital systems (e.g., Partners) using an affiliation list provided by the Connector.

This procedure should capture care given directly by the vast majority of Partners physicians. This includes specialists treating at the Partners hospital campuses, primary care physicians treating at Partners CHCs, and PCPs/specialists treating with the main Partners-owned medical groups (Mass General Physicians Organization, Brigham & Women's Physician Organization, Brigham Community Practices, Newton Wellesley-PHO, and North Shore Physicians Group). Statistics from Massachusetts' Registration of Provider Organization (RPO) dataset for 2015 suggest that over 90% of Partners-contracting physicians are part of these medical groups.[55] The measure will not capture physicians who are clinically affiliated with Partners but are independently owned or part of another health system so do not bill with Partners. My analysis of a clinical affiliation dataset for another project suggests that the vast majority (at least 80%) of Partners-affiliated physicians are also formally owned by Partners Healthcare System.[56]

A.3 Plan Choice and Cost Dataset Construction

The plan choice and cost dataset is described in Section 3.2. It includes a dataset of available plans, plan characteristics (including premium and network), and chosen options during fiscal 2008-2013. I also have data on fiscal 2014 choices, which I use for robustness checks on CeltiCare's network change (Section 4.2). However, I do not use it for the plan choice model or cost model estimation because I

[55] See RPO data publicly available at https://www.mass.gov/service-details/ma-rpo-data.

[56] The affiliation dataset comes from Massachusetts Health Quality Partners (see http://www.mhqp.org/resources-professionals/massachusetts-provider-directory-mpd/) but was purchased under a project-specific agreement so cannot be used for this paper without additional fees.

lack full claims data for 2014.

This dataset is constructed at the level of instances of enrollees making a plan choice. I start from the full enrollment dataset provided by the exchange, which includes one observation per member-month of enrollment with information on their enrolled plan and income group and demographics. I then limit this to the two instances where enrollees make a plan choice: (1) when an individual newly enrolls in CommCare (or re-enrolls after a gap), and (2) at annual open enrollment when current enrollees can switch plans. I make several exclusions from this sample for various reasons. Starting from a preliminary sample of 2,148,834 choice instances, I exclude 684 observations with missing/invalid income group or location data, 966 observations who enroll in a plan that is supposed to be unavailable based on their location, and 9,691 observations in the 200-300% of poverty income group who choose a lower-cost sharing option that was available only in 2007-08. Finally, I exclude 142,108 observations in the 0-100% of poverty group who were passively auto enrolled into a plan upon joining the exchange, since they do not make active choices that my plan choice model seeks to capture. The auto enrollment policy ended after 2009 so is not relevant for the main period of my study (see Shepard and Wagner (2020) for research studying this policy). The final sample includes 1,684,203 plan choice instances made by 624,443 unique enrollees. Summary statistics are shown in Table 5B.

Using administrative information from CommCare, I code the available plan choice set and the premiums and networks of each available plan. I define enrollee characteristics based on demographics on the enrollment file and information summarized from the linked claims data (e.g., medical conditions and HCC risk score). I use the available plan choice dataset along with enrollee characteristics to estimate the plan choice model described in Section 6.1. The sample counts in the plan choice model estimates (Table 10) differ slightly from those reported in Table 5B because the plan choice model drops 3.5% of instances where individuals have only a single plan available.

Table 5: Summary Statistics

Panel A: Hospitalization Dataset

Patient Characteristics

Variable		Mean
No. of Hospitalizations		70,094
Age		44.7
Male		48%
Emergency Department		65%
Principal	Mental Illness	14.9%
Diagnosis	Digestive	13.9%
	Circulatory	11.9%
	Injury / Poisoning	7.3%
	Respiratory	7.2%
	Cancer	6.8%
	Endocrine / Metabolic	6.3%
	Musculoskeletal	6.0%
	Pregnancy / Childbirth	5.4%
	All Other Diagnoses	20.4%

Chosen Hospital Statistics

Variable	Mean	Std. Dev.
Distance: Chosen Hosp. (miles)	12.7	15.1
All Hospitals (miles)	47.5	26.1
Hospital Category		
Academic Med. Ctr.	29%	---
Teaching Hospital	18%	---
All Others	53%	---
Partners Hospital	13%	---
Out-of-Network	8%	---
Past Use of Chosen Hospital (prior to this year)		
Any Use	43%	---
Inpatient Use	14%	---
Outpatient Use	42%	---
Total Cost to Insurer	$11,140	$14,017
Price (rel. to average)	1.019	0.274

Panel B: Plan Choice and Cost Dataset

Enrollee Characteristics

Variable	Mean	Std. Dev.
No. of Unique Enrollees	624,443	---
Age	39.9	14.0
Male	46.5%	---
Immigrant enrollee	5.6%	---
Income: <100% Poverty	46.8%	---
100-200% Poverty	39.4%	---
200-300% Poverty	13.7%	---
Past Use: Any Hospital	57.6%	---
Partners Hospitals	7.8%	---
Other 2012 Dropped Hosp.	5.3%	---
Risk Score: CommCare Score	1.001	0.924
HCC Risk Score	0.924	2.374
Choice Type: New Enrollee	29.5%	---
Re-Enrollee	13.7%	---
Current Enrollee	56.8%	---

Plan Statistics

Variable	Mean	Std. Dev.
No. of Choice Instances	1,684,203	---
Insurer Price (pre-subsidy)	$383.9	$69.6
Cons. Premium: Below Poverty	$0.0	$0.0
Above Poverty	$47.9	$46.1
Costs per Month: Total	$382.3	$1,484.5
Insurer Cost	$372.5	$1,478.6
Patient Cost Sharing	$9.7	$20.5
Hospital Network Utility	0.972	3.995
Share Covered Prev. Used Hosp.	0.740	0.420
Market Shares: BMC	35.7%	---
Network Health	34.4%	---
NHP	19.1%	---
CeltiCare	7.0%	---
Fallon	3.8%	---
Current Enr: Non-Switching	95.2%	---

NOTE: The table shows summary statistics for the hospitalization dataset (panel A) and the plan choice and cost dataset (panel B). These datasets are described Section 3.2. The hospitalization dataset is used to estimate the inpatient price model (Section 5.1 and Appendix C.2) and the hospital choice model (Appendix D.1). The plan choice and cost dataset is used to estimate the plan choice model (Section 6.1). The unit of observation for each sample is the "choice instance" – an inpatient hospitalization in panel A and an instance of making a plan choice in panel B. The latter occurs either when joining the exchange (new/re-enrollees) or during annual open enrollment when people can switch plans (current enrollees). Hospital network utility is a measure that enters the plan choice model and is described in Appendix D.2. The sample counts in Panel B differ slightly from the counts in the plan demand estimates in Table 10 because the latter excludes 3.5% of observations where there was only one available plan choice. These do not identify plan preferences but are included in the model analysis and simulations.

Appendix B: CommCare Premium and Network Variation

B.1 Prices, Subsidies, and Enrollee Premiums

My plan choice model (Section 6.1) is identified based on variation in plan prices and enrollee premiums. This appendix provides additional description on the pricing and subsidy institutions that lead to this variation. The starting point is pre-subsidy prices set by annual insurer bidding. Insurers submit sealed price bids to the regulator several months before the start of the plan year. The regulator then amalgamates these prices and applies subsidies, which determines enrollees premiums that apply at the start of the next plan fiscal year (which begins in July of the preceding calendar year; e.g., FY 2012 starts in July 2011). Prices and premiums are fixed for the remainder of the fiscal year. (Whenever not specified, years in the discussion below refer to fiscal years.)

Figure 9A shows average pre-subsidy prices in each CommCare fiscal year. (There are no points for 2008 because 2007 price bids were carried over to 2008 with an inflation update.) In 2007-2010, these prices represent enrollment-weighted averages across multiple pricing regions/cells. For 2007 and 2009, insurers could price separately by region, income group, and specified age-sex groups – with this more detailed pricing allowed because risk adjustment did not begin until 2010. In 2010, prices could be set at the region level (with five regions in the state). From 2011 on, insurers were required to set a single price for the whole state.

From pre-subsidy prices, subsidies were applied to generate post-subsidy "enrollee premiums." These vary substantially across income groups because of the application of different subsidies.[57] Average enrollee premiums are shown in Panel B of Figure 9, with separate averages for below-poverty and above-poverty income groups. The below-poverty group (black line) is fully subsidized, paying $0 for any available plan in all years. Above-poverty groups receive large subsidies but pay higher premiums on the margin for higher-price plans. The specific subsidies vary by income group in four bins: 100-150%, 150-200%, 200-250% and 250-300% of poverty. In general, subsidies are designed to be progressive both in levels and in differences. Lower-income groups pay less for all plans, and premium differences are narrower for lower- vs. higher-income groups.

For instance, consider premiums in 2012. Figure 9A shows the pre-subsidy prices, which vary by $87 per month across insurers – from a low of $360 for CeltiCare and Network Health to a high of $447 for BMC. For enrollee premiums, the below-poverty group pays $0 for any available plan. After subsidies, enrollees with incomes 100-150% of poverty pay premiums ranging from $0 for CeltiCare and Network Health up to $34 for BMC. Notice that subsidies substantially reduce both the level and difference in premiums between plans. Enrollees with incomes 150-200% of poverty pay premiums ranging from $39 for CeltiCare/Network Health up to $91 for BMC – a $52 difference. Enrollees with

[57] Two additional details are worth mentioning. First, while pre-subsidy prices could vary across age-sex groups in 2007-09, the exchange did not allow premiums to vary across these groups. Instead, they used a weighted-average composite bid across age groups to determine the pre-subsidy price for a given region x income group. Income-specific subsidies were then applied. Second, while insurers can only set prices at a region level (up to 2010) or statewide (2011+), sometimes post-subsidy premiums can vary across "service areas" within a region when the lowest-price plan is unavailable. When this occurs, the state adjusts subsidies so that the next cheapest plan has the targeted post-subsidy premium (e.g., $0 for 100-150% of FPL, $39 for 150-200% FPL). Plan availability can affect the level of plan premiums but does not affect premium differences across available plans. My demand model accounts for plan availability in the choice set definition.

incomes 200-250% of poverty pay premiums ranging from \$77 for CeltiCare/Network Health up to \$152 for BMC – a \$75 difference. Finally, enrollees with incomes 250-300% of poverty pay premiums ranging from \$116 for CeltiCare/Network Health up to \$197 for BMC – an \$81 difference. This example is representative of how subsidies affect both the level and difference in plan premiums in a progressive way.

B.2 Identifying Variation in Premiums

The subsidy schedule just described generates *within-plan* variation in premiums and premium changes that I use to identify premium coefficients in my plan demand model. Figure 10 gives an example for Network Health in the Boston region from 2010-2013. Panel A shows the levels of enrollee premiums by income group in each year. Panel B subtracts the premium of the cheapest available plan to show premium differences (or "relative premiums"), which are the key statistics for identifying price-sensitivity in a discrete choice model.[58]

The plot shows how changes in Network Health's (and its competitors') pre-subsidy prices (Figure 9A) translate through subsidies into *differential changes* across income groups in premiums for the same plan. For instance, Network Health's pre-subsidy price goes from being the lowest in 2010 to being second-lowest (after CeltiCare) in 2011. For enrollees, this results in a (post-subsidy) premium increase for all income groups 100-300% of the federal poverty level ("FPL") but no premium change for enrollees below 100% of FPL (who still pay \$0). Further, the *amount* of the premium increase varies from +\$10.38 for 100-150% FPL enrollees up to +\$29.85 for 250-300% of FPL enrollees. Figure 10 shows that across the four years shown, there is significant relative premium variation for Network Health, including both increases and decreases.

By comparing demand changes for the same plan across income groups – and especially relative to below-poverty enrollees who serve as a sort of "control group" for capturing unobserved quality – the model can infer a valid causal effect of premiums on demand. The difference-in-differences style logic and used of fixed effects is described in Section 6.1. Here is how it works for the example shown in Figure 10. First, the specification for plan utility (equation (11) in the text) includes plan-region-year dummies (ξ_{j,Reg_i,Yr_t}) that absorb any year-specific demand shock for Network Health in the Boston region. Thus, premium (and network) coefficients will be identified only by comparing demand *for the same plan across people* within a given region-year cell. Second, plan utility includes plan-region-income group dummies (ξ_{j,Reg_i,Inc_i}) that absorb any *persistent demand differences across income groups* for Network Health in Boston. The only remaining premium variation not captured by the fixed effects comes from the (within-plan, within-region) *differential changes* in premiums by income group.

The full plan demand model is estimated using all plans, regions, and income groups over the six years from 2008-2013. As noted in Appendix B.1, premiums are reset at the start of every fiscal year

[58] The cheapest premium is determined by the exchange's "price-linked" subsidies, which set subsidies so that the minimum post-subsidy price equals a target amount for each income group. In 2010-2012, the minimum premium for the five income groups shown are \$0, \$0, \$39, \$77, and \$116. In 2013, the min premium remains \$0 for the first two groups but rises to \$40, \$78, and \$118 for the next three groups.

Figure 9: CommCare Plan Prices and Enrollee Premiums

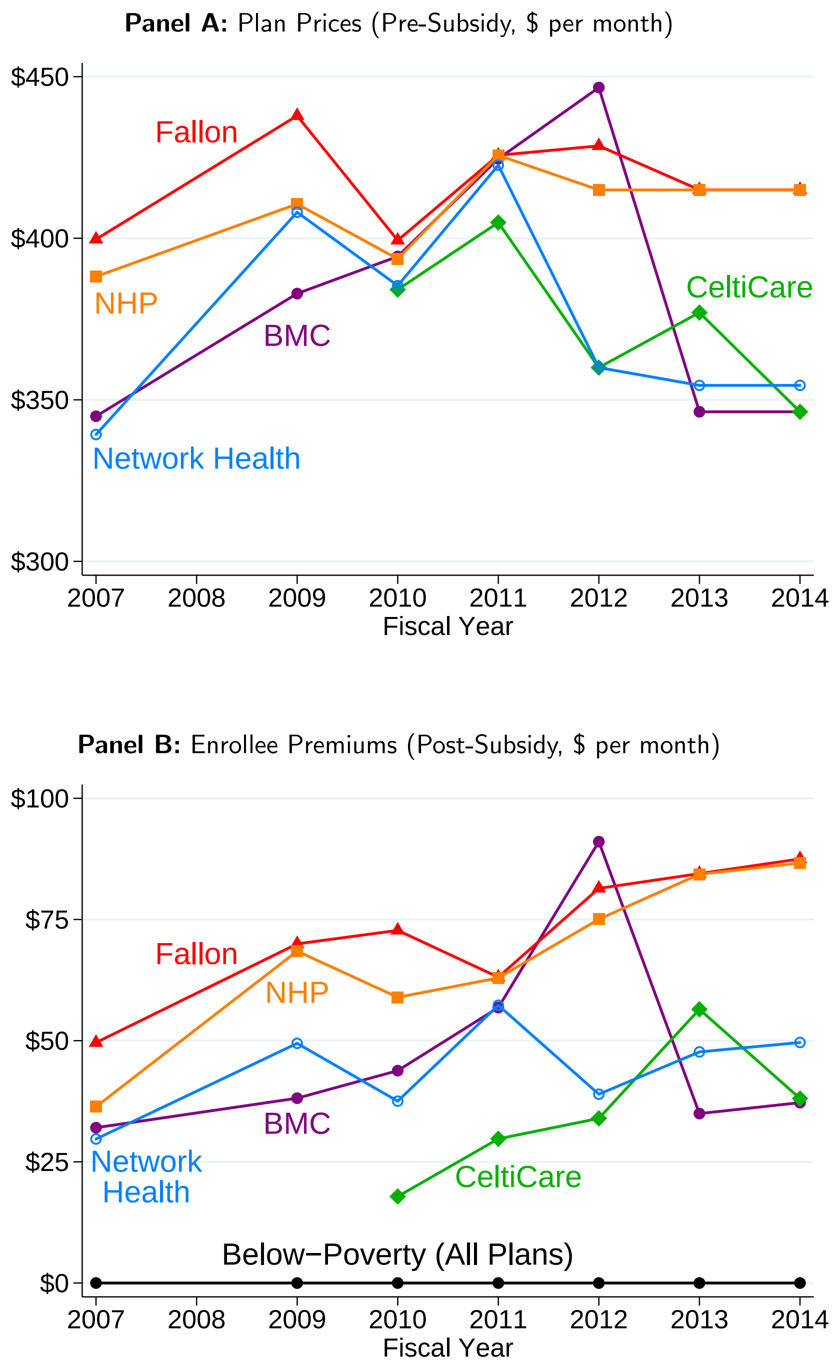

NOTE: The graphs show average pre-subsidy insurer prices (Panel A) and post-subsidy enrollee premiums (Panel B) for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying premiums and (in some years) prices vary by income group and region. The premiums in Panel B are shown separately for enrollees above-poverty (colored series) – who pay a subsidized amount related to the pre-subsidy price – and for below-poverty enrollees who are fully subsidized ($0 premium for all plans). I use the fact that subsidies imply different enrollee premiums for the same plans for identification of price sensitivity in my plan choice model.

Table 6: Distribution of Changes in Plan Relative Premiums

	Premium Decreases					Premium Increases				
	Share with	Distribution of Changes				Share with	Distribution of Changes			
	Decreases	Mean	Std Dev	Min	Max	Increases	Mean	Std Dev	Min	Max
All Years and Incomes	22%	-$31.0	$22.8	-$103.4	-$0.2	56%	$16.4	$15.9	$1.0	$103.4
By Income Group										
100-150% poverty	15%	-$22.5	$11.1	-$34.0	-$0.4	56%	$10.7	$8.6	$1.0	$35.1
150-200% poverty	22%	-$27.4	$20.5	-$57.1	-$0.2	55%	$15.3	$13.7	$2.0	$68.8
200-250% poverty	30%	-$42.2	$25.6	-$103.4	-$1.4	59%	$25.3	$20.8	$1.9	$103.4
250-300% poverty	35%	-$37.3	$29.7	-$103.4	-$0.6	55%	$28.2	$22.0	$1.6	$103.4
By Year										
2008-2009	18%	-$18.2	$17.9	-$53.7	-$2.6	30%	$43.1	$23.8	$5.6	$103.4
2009-2010	27%	-$34.6	$23.5	-$103.4	-$1.4	41%	$12.7	$10.0	$1.2	$60.9
2010-2011	5%	-$4.5	$5.4	-$25.9	-$0.2	86%	$11.1	$7.8	$1.3	$35.0
2011-2012	29%	-$18.4	$7.3	-$29.9	-$10.4	55%	$23.8	$13.7	$8.8	$81.0
2012-2013	27%	-$51.6	$18.4	-$81.0	-$1.0	70%	$8.2	$5.7	$1.0	$29.0

NOTE: The table shows statistics on the distribution of changes in (post-subsidy) enrollee premiums for each plan relative to the previous year. The underlying dataset includes one observation per plan x income group x service area x year cell (where service areas are the sub-region geographic level at which plan availability is determined) for the 2009-2013 period, excluding the income group 0-100% of poverty for whom all plans are $0 in all years. Statistics are calculated weighting by the number of enrollees in each cell. The variable of interest is the change in the plan's relative premium versus the previous year (for the same income group and service area). Relative premiums are defined as the plan's premium minus the cheapest available plan's premium; this nets out across-the-board shifts due to subsidy changes. The table shows the distribution separately for relative premium decreases and increases, along with the share of each. The remaining share of observations are cases with no change in the relative premium.

and are locked in for 12 months. Premiums for a given plan vary across income groups in all years and across regions prior to 2011. Table 6 shows the distribution of relative premium changes for a plan between adjacent years, separately for premium decreases and increases (following the presentation in Figure 7). The average relative premium decrease in the data is $31.0 per month, while the average premium increase is $16.4 per month. There is a substantial range of changes, with increases/decreases as large as $103 and as small as $1 or less. The table also shows how the distribution varies across income groups and years.

B.3 Hospital Networks

CommCare insurers have flexibility to set their covered hospital and medical provider network, subject to minimum network adequacy rules that were rarely binding. Figure 11 shows information on plans' share of hospitals covered (weighted by hospital beds), and Table 7 reports their coverage of the Partners Healthcare System hospitals. Through 2011, there were three broad-network plans: BMC HealthNet Plan, Neighborhood Health Plan (NHP), and Network Health. All of these covered about 80% of hospitals, and NHP and Network Health both covered most Partners hospitals. BMC did not cover Partners because it is owned by the rival Boston Medical Center hospital, but it otherwise has a broad network. Fallon is a regional plan based in central Massachusetts (and only available there in

Figure 10: Identifying Premium Variation Example: Network Health (Boston region), 2010-13

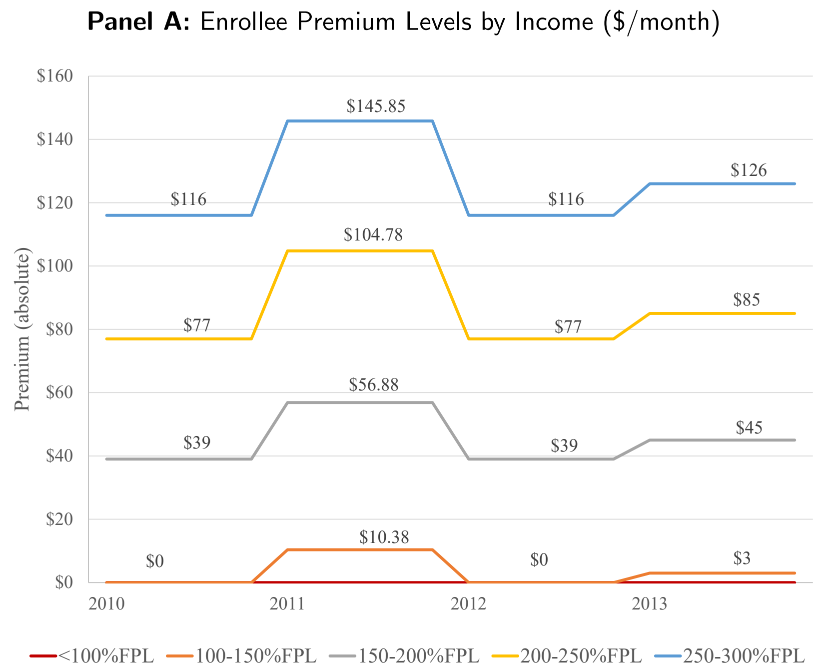

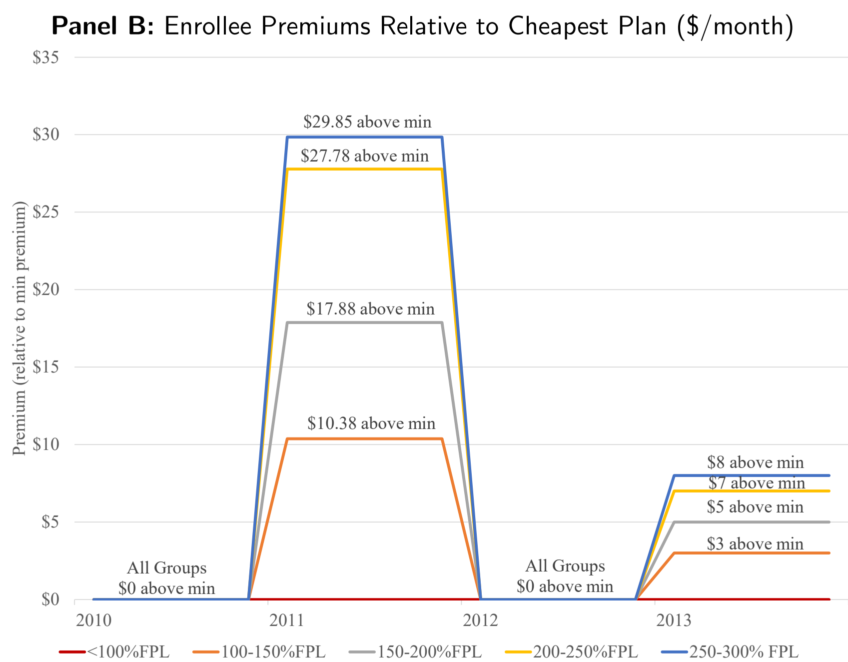

NOTE: The graphs shows the example of Network Health's (post-subsidy) enrollee premiums by income group over the 2010-2013 CommCare years. "FPL" refers to the federal poverty level. Pre-subsidy prices (and enrollee premiums) vary at the regional level in 2010, and the graph shows premiums specifically for the Boston region. Both are constant statewide in 2011-2013. Panel A shows the level of the premium for Network Health in dollars per month. Panel B shows the plan's "relative" premium, equal to the difference between its premium and the premium of the cheapest plan. The graph shows that different subsidies by income group translate a single pre-subsidy price into variation across income groups in the plan's post-subsidy relative premium.

Table 7: Coverage of Partners Hospitals by Exchange Plans

Plan	Hospitals	2009	2010	2011	2012	2013	*2014 (ACA)*
Boston Medical Center Plan (BMC)	MGH & Brigham	No	No	No	No	No	*No*
	Others	2/5	1/5	1/5	1/5	1/5	*0/5*
Network Health	MGH & Brigham	Yes	Yes	Yes	No	No	*No*
	Others	5/5	5/5	5/5	2/5	2/5	*0/5*
Neighborhood Health Plan (NHP)	MGH & Brigham	Yes	Yes	Yes	Yes	Yes	*Yes*
	Others	2/5	4/5	4/5	4/5	5/5	*5/5*
CeltiCare *(new in 2010)*	MGH & Brigham	---	Yes	Yes	Yes	Yes	*No*
	Others		3/5	3/5	3/5	3/5	*0/5*
Fallon *(mainly central MA)*	MGH & Brigham	No	No	No	No	No	*No*
	Others	0/5	0/5	0/5	1/5	0/5	*1/5*

NOTE: The table shows network coverage of the Partners hospitals by each CommCare plan over time. For each plan, the first line shows coverage of the two star academic hospitals – Mass. General Hospital (MGH) and Brigham & Women's Hospital – which are always bundled together. The next line shows how many of the five Partners community hospitals are covered in network.

later years), so it does not cover Partners hospitals and its statewide coverage is low.

CeltiCare is a new plan that enters the state in 2010 with a narrow network that covers less than half of hospitals but surprisingly, does cover Partners hospitals until 2014. It suffered from severe adverse selection after Network Health dropped Partners in 2012, and it subsequently decided to drop Partners in 2014. In testimony to the Mass. Health Policy Commission, CeltiCare's CEO wrote: "For the contract year 2012, Network Health Plan removed Partners hospital system and their PCPs from their covered network. As a result, the CeltiCare membership with a Partners PCP increased 57.9%. CeltiCare's members with a Partner's PCP were a higher acuity population and sought treatment at high cost facilities. . . . A mutual decision was made to terminate the relationship with BWH [Brigham & Women's] and MGH PCPs as of July 1, 2013." (Note that July 1, 2013, is the start of fiscal year 2014 for the purposes of the CommCare market.)

Network Health's dropping of Partners and several other hospitals in 2012 is evident in Figure 11 as the large fall in its hospital coverage share. It subsequently adds a few additional hospitals later in 2012-13, but it never restores coverage of Partners including after the ACA begins in 2014. Indeed, after its success in CommCare, it also dropped Partners in its (much larger) Medicaid managed care plan as of 2014. These changes left NHP as the only managed care plan that covers Partners in either Medicaid or the ACA "ConnectorCare" program that offers additional subsidies to low-income people in Massachusetts' ACA exchange.

Figure 11: Hospital Coverage in Massachusetts Exchange Plans

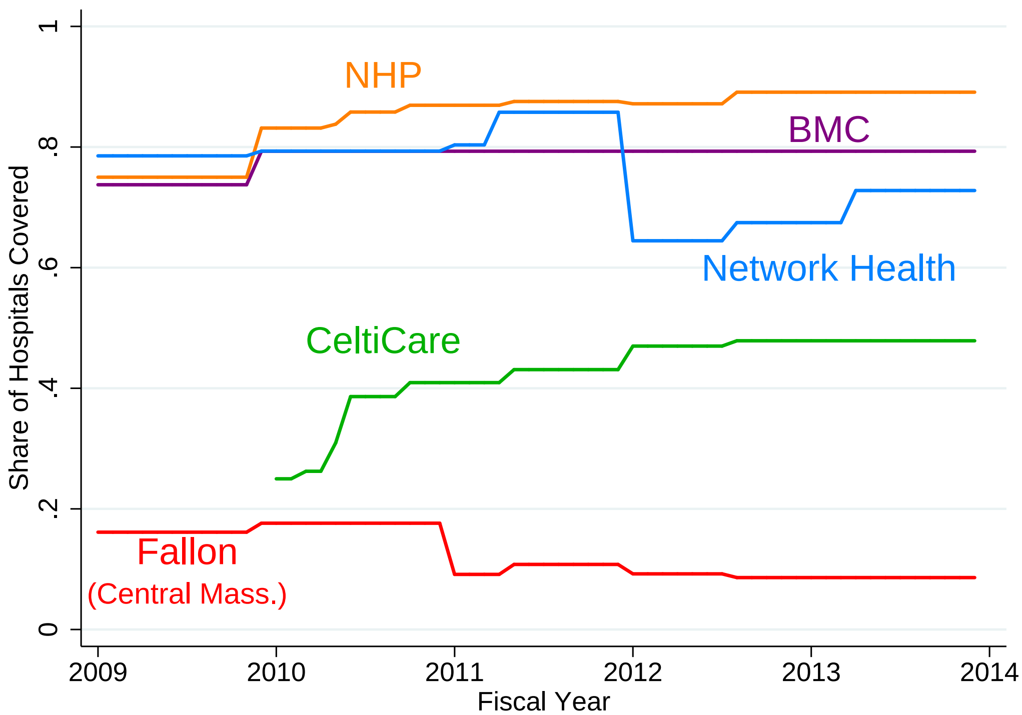

NOTE: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a statewide network.

Appendix C: Medical Spending Decomposition Details

This appendix describes the details of the method for decomposing medical spending, as summarized in Section 5.1. As noted in the text, the key steps involved are:

1. Defining the unit of medical services (s)
2. Estimating quantity for each medical service (Q_s) based on typical amounts paid for the service across all insurers and years
3. Calculating total quantity and average price for an enrollee

The following subsections describe how this is operationalized separately for outpatient and inpatient care. The final subsection reports some summary statistics on the share of cost variation accounted for by price versus quantity.

C.1 Outpatient Care

The most natural unit of service (s) for outpatient care are procedure codes, since the vast majority of care is paid for on a fee-for-service basis based on these. This definition, however, means that I exclude outpatient care that is paid for via other methods like capitation. In practice, non-FFS payments are not very common in the claims data.[59] I also exclude outpatient emergency department care to avoid double-counting, since these are included in the inpatient costs (see Appendix C.2 below) when there is an inpatient admission. Therefore, my outpatient cost decomposition reflects non-emergency department outpatient care.

I define a unit of service, s, based on HCPCS procedure codes (as used by Medicare and most private insurers, including CommCare) interacted with the type of bill/provider. HCPCS codes are detailed service units; an example code is 99213, a 15-minute physician office visit with an established patient. The type of bill/provider captures the distinction between bills for facility costs vs. professional services, as well as high-level provider categories (e.g., medical, behavioral health, and dental care) for which a given procedure may mean something slightly different. Following Medicare rules, a procedure delivered in a "facility" (e.g., a hospital or nursing home) is billed in two parts, with one payment for facility costs and one payment to the physician for professional services. I treat these bills as separate "services" and use each one's average price to calculate price-standardized utilization.

Given this definition of s, I define quantity Q_s as the mean insurer-paid amount ($Paid_{a_{it},s}$ in the notation of Section 5.1) for the service across all insurers and years of the claims data. Price is defined as the residual multiplicative factor that accounts for observed spending: $P_{a_{it},s} \equiv Paid_{a_{it},s}/Q_s$. This ensures that price measures are centered around 1.0. It also means that total quantity is a form of price-standardized utilization, which adds up services used valued at constant prices across insurers and years.

[59] Public reports indicate very little capitation payment by CommCare insurers. This is consistent with my analysis of the claims data, for which just 0.4% of claim lines for outpatient care (representing 0.6% of spending) have flags indicating capitation contracts. I exclude these claims from the outpatient cost decomposition.

With these definitions (and following Section 5.1), total quantity of outpatient care for an enrollee equals $Q^{OP}_{i,t} = \sum_{a_{it} \in A^{OP}_{it}} Q_{s(a_{it})}$, and average price equals $P^{OP}_{i,t} \equiv \frac{C^{OP}_{i,t}}{Q^{OP}_{i,t}}$. Finally, I can divide by its mean in a sample ($\overline{Q}$) to calculate relative quantity $Q^{rel,OP}_{i,t} \equiv Q^{OP}_{i,t}/\overline{Q}$.

C.2 Inpatient Care

For inpatient care, the most natural service unit is the diagnosis-related group (DRG), which is the standard measure used in hospital price analyses (e.g., Cooper et al., 2019) and is the method of payment for about 90% of hospitalizations in my data. Nonetheless, because not all admissions are DRG-paid and because even DRG payment allows exceptions due to "outlier adjustments," I estimate a pricing model that allows quantity to vary within a DRG or diagnosis based on other patient severity observables. Essentially, this method defines the quantity associated with each hospital admission in a continuous way based on a projection of spending onto DRG/diagnosis categories and other patient observables.

Consider a particular admission a – for enrollee i in plan j in year t for DRG (or diagnosis) d at hospital h.[60] I regress log insurer payments ($\log(Paid_{a,i,j,t,d,h})$) on insurer-hospital dummies $\alpha_{h,j,N}$ that can vary with the network status ($N \in \{0,1\}$), year dummies (β_t), DRG/diagnosis fixed effects (γ_d), and patient severity factors ($Z_{a,i,t}$) comprised of gender x age groups (in 5-year bins), income groups, and Elixhauser comorbidities:[61]

$$\log\left(Paid_{a,i,j,t,d,h}\right) = \alpha_{h,j,N} + \beta_t + \gamma_d + Z_{a,i,t}\delta + u_{a,i,j,d,t} \tag{13}$$

Using estimates of (13), I define the quantity unit as the component of payment arising from DRG/diagnosis, severity, and the residual, converting the estimate to spending levels:

$$\widetilde{Q}_{a,i,t} \equiv \exp\left(\hat{\gamma}_d + Z_{a,i,t}\hat{\delta} + \hat{u}_{a,i,j,d,t}\right) \tag{14}$$

The residual ($\hat{u}$) seems most natural to treat as quantity, since it likely reflects outlier adjustments and unmeasured add-on services. The remainder of (13) is defined as price:

$$\widetilde{P}_{a,i,j,h,t} \equiv \exp\left(\hat{\alpha}_{h,j,N(h,t)} + \hat{\beta}_t\right) \tag{15}$$

where I rescale the (non-identified) constant multiplier between price and quantity so that $\widetilde{P}_{j,h,t}$ has mean of 1.0 across the full sample (which means that $\widetilde{Q}$ is denominated in dollars). Given these

[60] When the DRG is unavailable, I use the single-level Clinical Classification Software (CCS) category of the principal diagnosis. CCS codes are a categorization defined by the U.S. Agency for Healthcare Research and Quality. As an alternative, I considered using DRG grouper software to impute the DRG for admissions where it is not listed. I found, however, that the claims data often did not include all necessary information to impute DRGs, making this method unreliable. The main missing information was ICD-9 procedure codes for the inpatient facility bill, which is required by Medicare DRG grouper software.

[61] This regression specification is quite similar to that of Cooper et al. (2019). To avoid over-fitting, I pool $\alpha_{h,j,Netw}$ cells with fewer than 11 observations into an "other hospitals" group, still separately by insurer and network status. This pooling only applies to about 0.5% of admissions – primarily for out-of-network care and small hospitals, and I ensure it does not affect the star hospitals.

definitions of price and quantity, I apply the decomposition framework to define inpatient quantity for i in year t as: $Q_{i,t}^{IP} \equiv \sum_{a \in Admit(i,t)} \widetilde{Q}_{a,i,t}$. Inpatient price is $P_{i,t}^{IP} \equiv C_{i,t}^{IP}/Q_{i,t}^{IP}$. I also define relative quantity as above.

Inpatient and outpatient care estimates can be analyzed separately or combined to form a decomposition for total costs in the sample. If combined, total quantity equals the sum of the two:

$$Q_{i,t}^{Tot} \equiv Q_{i,t}^{IP} + Q_{i,t}^{OP}$$

and relative quantity can also be defined ($Q_{i,t}^{rel} \equiv Q_{i,t}/\overline{Q}$) by taking out the sample mean $\overline{Q}$. Price is defined as the remaining factor needed to account for costs (which as noted in Section 5.1 equals a weighted average of service-level prices):

$$P_{i,t}^{Tot} = \frac{C_{i,t}^{IP} + C_{i,t}^{OP}}{Q_{i,t}^{Tot}}$$

C.3 Summary of Estimates

Appendix Table 8 shows a summary of the decomposition. Panel A shows statistics about the mean and standard deviation of medical costs and the quantity and price decomposition estimates. Panel B shows the relationship of quantity and price to CommCare's medical risk score. For both analyses, the unit of analysis is the enrollee-year (reflecting the insurance contract period), and the sample is limited to 2011-2013, the years risk scores are available. All results are similar if I instead restrict the analysis to Network Health in 2011 (the key plan-year for the selection analysis).

Panel A shows that there is substantial cost variation across enrollees, with both quantity and price contributing. For total costs covered by the decomposition (column (1)), its mean is \$228.5 per month (which, recall, is 61% of overall costs). Its standard deviation of \$780.5 is more than three times as large, reflecting the skewed nature of medical spending. Most of this variation comes from quantity, whose coefficient of variation is 3.15. But price also varies meaningfully, with a standard deviation of 34% across enrollees. (Interestingly, price and quantity are largely orthogonal, with a correlation of -0.02.) The same basic patterns hold separately for outpatient and inpatient costs in columns (2)-(3).

Panel B shows the relationship of this quantity/price variation to the enrollee risk scores CommCare used for risk adjustment, using simple regressions of quantity/price on risk score and a constant. This relationship is important for selection incentives: the better risk scores capture predictable cost variation, the more likely they will neutralize selection incentives. The table shows that while risk scores strongly predict quantity of care – with a regression coefficient of 1.059 (s.e. = 0.017) – they hardly predict price variation at all (coeff. = 0.002, s.e. = 0.001). Similarly, the R^2 is about 10% for quantity versus <0.01% for price. This pattern is consistent across total, outpatient, and inpatient costs: in all cases, risk score strongly predicts quantity but barely predicts price.[62]

Overall, Table 8 suggests that while utilization is the main driver of cost heterogeneity, the price

[62] Although not shown, this pattern also holds if I use the concurrent HCC risk score used by the ACA instead of CommCare's retrospective score. Indeed, the HCC score does an even better job capturing quantity variation ($R^2 = 26\%$), while R^2 for price remains $< 0.01\%$.

Table 8: Price vs. Quantity Medical Cost Decomposition

Variable	Statistic	Total Costs (1)	Outpatient Costs (2)	Inpatient Costs (3)
Panel A: Cost Decomposition Summary				
Costs in Decomp.	Mean	$228.5	$163.6	$65.0
($ per month)	*[S.D.]*	*[$780.5]*	*[$388.7]*	*[$610.1]*
Relative Quantity	Mean	1.00	1.00	1.00
	[S.D.]	*[3.15]*	*[2.39]*	*[8.52]*
Relative Price	Mean	1.02	1.02	1.00
	[S.D.]	*[0.34]*	*[0.34]*	*[0.26]*
Panel B: Regression of Quantity/Price on Risk Score				
Relative Quantity	Regr. Coeff	1.059	0.846	1.620
	(s.e.)	(0.017)	(0.015)	(0.045)
	[R^2]	*[9.6%]*	*[10.6%]*	*[3.1%]*
Relative Price	Regr. Coeff	0.002	0.002	0.008
	(s.e.)	(0.001)	(0.001)	(0.001)
	[R^2]	*[0.0%]*	*[0.0%]*	*[0.4%]*

NOTE: The table shows a summary of the decomposition of medical costs into price versus quantity. Panel A shows means and standard deviations across enrollees for costs included in the decomposition (in $ per member-month) and for the relative quantity and price estimates (see text for their definitions). Panel B shows estimates of regressions of quantity/price (y-variable) on an enrollee's CommCare risk score. For both panels, the columns show results separately for (1) total costs in the decomposition, (2) outpatient costs, and (3) inpatient costs. Observations are at the enrollee-year level (with outcomes averaged to per-month values) and are weighted by number of months a person is enrolled during the year. The sample is limited to fiscal years 2011-2013, the years for which risk scores are available.

dimension of costs – reflecting enrollees' use of higher-price providers – is also relevant. Moreover, the price dimension is not well captured by risk adjustment, consistent with it being driven by a different source of heterogeneity than the sickness measures that enter risk adjustment. This suggests that both (residual) quantity and price variation may be important for insurer selection incentives.

Appendix D: Structural Model and Estimation Details

D.1 Hospital Choice Model

I use the inpatient hospitalization dataset (see Appendix A.1) to estimate a multinomial logit choice model. I distinguish patients' utility for different hospitals from the barriers their plan's network creates. The utility of patient i with diagnosis d for hospital h at time t is:

$$U_{i,d,t,h}^{Hosp} = \underbrace{\gamma_1\left(Z_{i,d,t}\right) \cdot Dist_{i,h}}_{\text{Distance}} + \underbrace{\gamma_2\left(Z_{i,d,t}\right) \cdot X_h + \gamma_3 \cdot PastPatient_{i,h,t}}_{\text{Hospital characteristics x Patient observables}} + \underbrace{\eta_h}_{\text{Hospital dummy}} + \underbrace{\epsilon_{i,d,t,h}}_{\text{Logit error}} \tag{16}$$

The function governing patient choices (and entering the logit equation) equals this utility minus a hassle cost of going out of network:

$$u_{i,j,d,t,h}^{Hosp} = U_{i,d,t,h}^{Hosp} - \kappa_j\left(Z_{i,t}\right) \cdot 1\left\{h \notin N_{j,t}\right\} \tag{17}$$

The specification in (16) is similar to past work (e.g., Town and Vistnes, 2001; Gaynor and Vogt, 2003; Ho, 2006). While this past work (if it measures networks at all) simply excludes out-of-network hospitals from the choice set, I include these hospitals and instead estimate an out-of-network hassle cost $\kappa_j\left(Z_{i,t}\right)$, which can vary by insurer and patient severity and emergency status. I choose this approach because of the observation that a non-trivial share of patients (about 8%) use out of network hospitals, both for emergencies and non-emergencies. This can occur when the insurer gives prior authorization to go out of network, a barrier that is reasonably represented as a hassle cost. Notice that my approach is a generalization of the standard practice of excluding out-of-network hospitals from the choice set; my model's predictions converge to the standard approach as $\kappa \to \infty$.

In addition to hospital dummies, the utility covariates in (16) include patient travel distance and patient observables interacted with hospital characteristics to allow patient preferences and substitution patterns to differ. The distance variables include distance (in miles) and distance-squared (with separate coefficients for patients living in each of five regions of the state) and distance interacted with patient age, gender, income group, emergency status, and severity (the $\widetilde{Q}_{a,i,t}$ metric from the price decomposition; see equation (14)). The patient observable x hospital characteristics variables are: (1) patient diagnosis category (using the top-level CCS category) interacted with hospital's service offerings (e.g., cancer patient x hospital has oncology services); (2) hospital academic type (top academic medical center, teaching hospital, community hospital) interacted with patient severity, diagnosis category, and whether the patient is a past Partners patient; and (3) whether patient i has previously used hospital h or its doctors (separate dummies for inpatient and outpatient care) prior to the current plan year (and at least 30 days prior to the admission, to avoid any mechanical relationship).

Including past provider use variables differs from past work, which has often not had panel data or outpatient claims to measure it. Including past use allows me to capture relationships between patients and a hospital's physicians, which is a key source of heterogeneity in hospital choices. However, this coefficient's interpretation is complicated because it picks up both state dependence and heterogeneity (see analogous discussion in Section (4.3)). To deal with this issue, I assume is that these relationships are fixed in the short run – e.g., the one-year horizon in my counterfactuals – so past use variables are held fixed in all simulations. Of course, it would be nice to model the process through which these patient-provider relationships form. But doing so would introduce complicated dynamics into an already complex model. Instead, I treat these relationships as exogenous, which is sensible in the short run (but less ideal over longer horizons).

Estimates Because all covariates are observed, I estimate the model by maximum likelihood. Table 9 shows the results. Consistent with previous papers' estimates, patients dislike traveling to more distant hospitals, with each extra mile of distance reducing a hospital's share by 7.6% on average. The model estimates a sizeable hassle cost for out-of-network hospitals that reduces their shares by 58% on average. Two sets of coefficients have implications for the main selection findings of the paper. First, teaching hospitals and academic medical centers (AMCs) tend to attract sicker patients, both measured by patient severity and by particular diagnoses (e.g., cancer). Moreover, AMCs and teaching hospitals are particularly attractive to past Partners patients. Second, past care use is a very strong predictor of future hospital choices. Patients choose a hospital where they have a relationship about 40% of the time, about twice as high as would be expected based on other covariates.

D.2 Hospital Network Utility

To generate a measure of network utility for plan demand, I follow the method of Capps et al. (2003). Consider a consumer i who is deciding among various plans j (with networks $N_{j,t}$) at time t. I define network utility of each plan based on the expected utility metric from the hospital demand system. Conditional on needing to be hospitalized for diagnosis d with emergency status $e \in \{0,1\}$, at time t, a consumer's utility of access to network $N_{j,t}$ in plan j is:

$$\begin{aligned} EU_{i,d,e,t,j}\left(N_{j,t}\right) &= E\left[\max\left\{\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right) + \varepsilon_{i,d,e,t,h}\right\}\right] \\ &= \log\left(\sum_{h} \exp\left(\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right)\right)\right) \end{aligned} \tag{18}$$

where $\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right)$ is the utility function from (17) excluding the logit error term. (Note that I explicitly include emergency status e in the subscripts here; in equation (16) it was implicitly part of $Z_{i,d,t}$.) Many covariates that enter hospital utility are known at the time of plan choice (e.g., distance, past patient status, and demographics). However, other variables are not realized until later: notably diagnosis, emergency status, and severity. I assume that consumers have expectations over these variables based on observed patterns in the data. Consumers have expectations for their hospital use frequency for each diagnosis d and emergency status $e \in \{0,1\}$ over the coming year, which I denote $freq_{i,d,e,t}$. I estimate these frequencies using a Poisson regression of the number of hospitalizations in the data (for a given $\{d,e\}$ combination) on age-sex and income groups.[63] I use the predicted values from these regressions for $freq_{i,d,e,t}$. For patient severity, I use the average observed severity in the

[63] I choose not to use diagnoses in this regression because past diagnoses are unavailable for new enrollees.

Table 9: Hospital Choice Model Estimates

Variable	Coeff.	Std. Error
Distance to Hospital (miles):		
Distance (base coeff.: Boston)	-0.2320	(0.0052)
x Region = Central Mass.	0.0889	(0.0057)
x Region = Northern Mass.	0.0561	(0.0058)
x Region = Southern Mass.	0.1030	(0.0052)
x Region = Western Mass.	0.1452	(0.0058)
Distance^2 (avg. coeff.)	0.0012	(0.00002)
Distance x 1{Income > Poverty} (avg.)	-0.0080	(0.0009)
x Age / 10	-0.0031	(0.0003)
x Male	0.0063	(0.0009)
x Admission Severity	0.0021	(0.0006)
x Emergency	-0.0203	(0.0009)
Past Patient of this Hospital		
Inpatient Care	0.9958	(0.0390)
Outpatient Care	1.8195	(0.0200)
Hospital x Patient Characteristics		
Academic Med. Ctr. x Severity	0.4300	(0.0377)
Teaching Hospital x Severity	0.2261	(0.0336)
AMC x Past Partners Patient	0.3224	(0.0569)
Teaching x Past Partners Patient	0.3508	(0.0647)
AMC/Teaching x Diagnoses	Yes	
Selected Coeffs: AMC x Cancer	1.3257	(0.0666)
AMC x Injury	1.0210	(0.0953)
AMC x Musculosk.	0.4308	(0.0903)
AMC x Mental	-1.4726	(0.0626)
Diagnosis x Hospital Specialty Services	Yes	
Hospital Dummy Variables	Yes	
Out-of-Network Disutility		
Out-of-Network x Plan = BMC	-1.8590	(0.0517)
x Plan = CeltiCare	-2.3100	(0.0732)
x Plan = Fallon	-1.8027	(0.0748)
x Plan = NHP	-0.9391	(0.0652)
x Plan = Network	-1.8405	(0.0495)
Out-of-Network x Emergency	0.9084	(0.0433)
Model Stats: Number of Admissions	70,094	
Number of Individuals	47,958	
Pseudo-R^2	0.578	

NOTE: The table shows estimates for the multinomial logit hospital choice model. The coefficients shown are interpretable as entering the utility function describing hospital choice. Past use variables are dummies for whether a patient has previously used each specific hospital (before the current plan year and at least 30 days before the current admission). Severity is an estimated summary measure ($\widetilde{Q}_{a,i,t}$) from the inpatient price model described in Appendix C; it is standardized (mean 0, SD 1) before entering as a covariate in this model. In addition to the variables shown, the model includes: distance interacted with detailed income group (0-100% poverty and by 50% of poverty from 100-300%); distance-squared interacted with region; interactions between academic medical center (AMC) and teaching hospital status and diagnoses; and seven diagnosis x hospital specialty service interactions (cancer x oncology services; cardiovascular diagnosis x cath lab, x interventional cardiology, and x heart surgery services; pregnancy x obstetrics services and x NICU; musculoskeletal diagnosis x arthritis services; and injury diagnosis x level 1 trauma center).

hospitalization data for the $\{d, e\}$ and age-sex group cell.

Given these expectations, the *ex-ante* expected network utility is:

$$NetworkUtil_{i,j,t}\left(N_{j,t}\right) = \sum_{d,e} freq_{i,d,e,t} \cdot EU_{i,d,e,t,j}\left(N_{j,t}\right) \tag{19}$$

The network utility in (19) is what I include in plan demand. Because network utility does not have natural units, I normalize it so that 1.0 is the average decrease in utility for Boston-region residents when Network Health dropped Partners in 2012.

D.3 Plan Choice Model Details

The plan choice model is described in Section 6.1. Table 10 below shows a summary of estimates, and Table 11 lists the full set of coefficients on plan attributes (premium, network value, and inertia) that enter the model, including interaction terms with enrollee observables. The plan choice model also includes a large number of plan dummy variables and interactions (251 coefficients including all interactions), which are not reported (but will be available in data output in the replication packet).

Table 10: Insurance Plan Choice Model Estimates

Variable	(1) New/Re-Enr. Only Coeff.	(1) New/Re-Enr. Only Std. Error	(2) All Enrollees Coeff.	(2) All Enrollees Std. Error
Enrollee Premium (per $10/month): Avg. Coeff.	**-0.454**	**(0.004)**	**-0.506**	**(0.003)**
Base Coeffs by Income: 100-150% poverty	**-0.734**	**(0.010)**	**-0.774**	**(0.008)**
150-200% povery	**-0.506**	**(0.009)**	**-0.564**	**(0.008)**
200-250% poverty	**-0.415**	**(0.008)**	**-0.451**	**(0.007)**
250-300% poverty	**-0.392**	**(0.009)**	**-0.424**	**(0.007)**
x High Risk Score (>80th pctile)	0.084	(0.009)	0.089	(0.008)
x Any Chronic Illness	0.018	(0.003)	0.018	(0.003)
x Cancer	0.041	(0.005)	0.037	(0.004)
x Age ≥45 years	0.111	(0.011)	0.094	(0.010)
Provider Network				
Network Utility (avg. coeff.)	**0.506**	**(0.005)**	**0.463**	**(0.005)**
x Income >100% poverty.	0.097	(0.008)	0.059	(0.007)
x High Risk Score (>80th pctile)	-0.252	(0.014)	-0.239	(0.013)
x Any Chronic Illness	0.135	(0.006)	0.129	(0.005)
x Cancer	0.040	(0.011)	0.033	(0.010)
Share Prev Used Hosp. Covered (avg. coeff.)	**0.249**	**(0.013)**	**0.291**	**(0.012)**
x Income >100% poverty.	0.217	(0.026)	-0.011	(0.022)
x High Risk Score (>80th pctile)	0.277	(0.044)	0.262	(0.037)
x Any Chronic Illness	0.203	(0.027)	0.164	(0.022)
x Cancer	0.129	(0.053)	0.188	(0.041)
x Prev. Used Partners Hospitals	0.625	(0.023)	0.982	(0.021)
Inertia: Current Plan Dummy (avg. coeff.)	---		**4.413**	**(0.007)**
x Income >100% poverty.	---		-1.059	(0.013)
x High Risk Score (>80th pctile)	---		-0.136	(0.032)
x Any Chronic Illness	---		-0.153	(0.013)
x Age ≥45 years	---		-0.079	(0.020)
Avg. Plan Dummies: BMC HealthNet	*(normalized = 0)*		*(normalized = 0)*	
CeltiCare	-1.055	(0.029)	-1.082	(0.025)
Fallon	-0.049	(0.040)	0.058	(0.034)
Neighborhood Health Plan (NHP)	-0.090	(0.016)	-0.037	(0.015)
Network Health	-0.001	(0.013)	-0.119	(0.012)
Model Stats: Pseudo-R^2	0.181		0.575	
No. Choice Instances	690,365		1,613,003	
No. Unique Enrollees	526,665		611,070	

NOTE: This table shows estimates for the multinomial logit plan choice model described in Section 6.1. Column (1) includes just new and re-enrollees who make active choices (so do not have inertia terms). Column (2) shows the main model that includes all enrollees, with inertia variables for current enrollees. Premium is the amount paid by consumers after subsidies, in $10 per month; this varies by about $20-60 across plans. Network utility is the consumer-specific expected utility measure for a plan's hospital network, defined in Appendix D.2. Share previously used hospitals covered is the share of an enrollee's previously used hospitals that a plan covers, with a separate interaction for the star Partners hospitals. For most covariates, I report the average coefficient across all enrollees, as well as selected interactions terms with consumer observables. The model allows for more interactions than those shown. For premium and inertia, it includes interactions with: (1) income groups, (2) risk score quantiles (quintiles with a separate category for the 95-100 percentiles), (3) diagnosis indicators (chronic disease, cancer), (4) demographics (5-year age-sex groups and immigrant status). The provider network measures are interacted with all of these except demographics. Plan dummies are interacted with region-year dummies, region-income dummies, and risk score quantiles and demographics.

Table 11: Plan Choice Model: Full List of Coefficient Estimates

Variable	Coeff.	Std. Error
Enrollee Premium (per $10/month)		
Base Coeffs: 100-150% FPL	**-0.774**	**(0.008)**
150-200% FPL	**-0.564**	**(0.008)**
200-250% FPL	**-0.451**	**(0.007)**
250-300% FPL	**-0.424**	**(0.007)**
x Risk Score: 20-40th pctile	0.014	(0.008)
40-60th pctile	0.027	(0.009)
60-80th pctile	0.038	(0.008)
80-95th pctile	0.075	(0.008)
95-100th pctile	0.129	(0.009)
x Any Chronic Illness	0.018	(0.003)
x Cancer	0.037	(0.004)
x Age-Sex Grp: Male 19-24	(omitted)	
Male 25-29	0.014	(0.009)
Male 30-34	0.033	(0.009)
Male 35-39	0.060	(0.012)
Male 40-44	0.066	(0.011)
Male 45-49	0.079	(0.011)
Male 50-54	0.088	(0.011)
Male 55-59	0.084	(0.011)
Male 60+	0.099	(0.011)
Female 19-2	-0.008	(0.009)
Female 25-2	0.044	(0.011)
Female 30-3	0.070	(0.011)
Female 35-3	0.089	(0.012)
Female 40-4	0.095	(0.011)
Female 45-4	0.085	(0.011)
Female 50-5	0.089	(0.011)
Female 55-5	0.094	(0.011)
Female 60+	0.130	(0.011)
x Immigrant enrollee	-0.259	(0.010)

Variable	Coeff.	Std. Error
Network Utility: Base Coeff.	**0.551**	**(0.012)**
x Income: 100-150% FPL	0.086	(0.009)
150-200% FPL	0.053	(0.010)
200-250% FPL	0.019	(0.011)
250-300% FPL	0.029	(0.015)
x Risk Score: 20-40th pctile	-0.244	(0.012)
40-60th pctile	-0.243	(0.012)
60-80th pctile	-0.303	(0.012)
80-95th pctile	-0.243	(0.014)
95-100th pctile	-0.229	(0.018)
x Any Chronic Illness	0.129	(0.005)
x Cancer	0.033	(0.010)
Share Prev. Used Hosp Covere	**0.052**	**(0.029)**
x Income: 100-150% FPL	-0.162	(0.027)
150-200% FPL	0.127	(0.029)
200-250% FPL	0.026	(0.034)
250-300% FPL	0.117	(0.044)
x Risk Score: 20-40th pctile	0.121	(0.032)
40-60th pctile	0.080	(0.034)
60-80th pctile	0.171	(0.035)
80-95th pctile	0.241	(0.039)
95-100th pctile	0.323	(0.059)
x Any Chronic Illness	0.164	(0.022)
x Cancer	0.188	(0.041)
x Prev. Used Partners Covere	0.735	(0.058)
x Income: 100-150% FPL	0.051	(0.053)
150-200% FPL	-0.358	(0.054)
200-250% FPL	-0.245	(0.061)
250-300% FPL	-0.357	(0.075)
x Risk Score: 20-40th pct	0.272	(0.066)
40-60th pcti	0.482	(0.068)
60-80th pcti	0.685	(0.070)
80-95th pcti	0.466	(0.073)
95-100th pc	0.526	(0.093)
x Any Chronic Illness	-0.086	(0.044)
x Cancer	0.185	(0.062)

Variable	Coeff.	Std. Error
Inertia (Current Plan Dummy)	**5.435**	**(0.027)**
x Income: 100-150% FPL	-0.872	(0.016)
150-200% FPI	-1.023	(0.017)
200-250% FPI	-1.419	(0.018)
250-300% FPI	-1.372	(0.023)
x Risk Score: 20-40th pctile	-0.107	(0.033)
40-60th pctile	-0.113	(0.036)
60-80th pctile	-0.184	(0.033)
80-95th pctile	-0.118	(0.032)
95-100th pctile	-0.188	(0.040)
x Any Chronic Illness	-0.153	(0.013)
x Cancer	-0.079	(0.020)
x Age-Sex Grp: Male 19-24	(omitted)	
Male 25-29	-0.072	(0.037)
Male 30-34	-0.125	(0.040)
Male 35-39	-0.129	(0.048)
Male 40-44	-0.149	(0.047)
Male 45-49	-0.154	(0.046)
Male 50-54	-0.206	(0.045)
Male 55-59	-0.249	(0.046)
Male 60+	-0.251	(0.046)
Female 19-2	-0.136	(0.034)
Female 25-2	-0.258	(0.044)
Female 30-3	-0.350	(0.046)
Female 35-3	-0.301	(0.049)
Female 40-4	-0.278	(0.047)
Female 45-4	-0.336	(0.045)
Female 50-5	-0.357	(0.044)
Female 55-5	-0.430	(0.043)
Female 60+	-0.344	(0.043)
x Immigrant enrollee	-0.408	(0.025)

D.4 Cost Model Estimates

The insurer cost model is described in Section 6.2 and is based on the reduced form analysis in Section 5.3. Table 12 shows estimates of the key piece of the model: how costs change at the enrollee level due to the narrower network adopted by Network Health in 2012. The estimating equation is:

$$E\left(C_{i,j,t}\right)=\exp\left(\alpha_i+\beta_t\left(Z_i\right)+\gamma\left(Z_i\right)\cdot 1\left\{j=NH,t\geq 2012\right\}\right) \tag{20}$$

where $C_{i,j,t}$ is insurer cost on individual i at time t, α_i is an enrollee fixed effect (which is divided out and not estimated), $\beta_t(.)$ are time fixed effects that capture trends for the control group, and Z_i are enrollee characteristics on which time trends and causal effects may vary. Regression (20) is estimated by maximum likelihood (using "xtpoisson, fe" in Stata), with cluster-robust standard errors at the i level. The coefficients of interest are $\gamma\left(Z_i\right)$, which capture the differential cost change for Network Health stayers in 2012.

Table 12 shows the estimates of $\hat{\gamma}\left(Z_i\right)$, the key coefficients of interest. Recall that the implied (multiplicative) effect on costs equals $dC_i=\exp\left(\hat{\gamma}\left(Z_i\right)\right)$, and the percent change is dC_i-1. Columns (1)-(3) report models with increasing flexibility in the Z_i with which γ is allowed to vary. Column (3) is the full model that is used for the final cost analysis in Sections 6.3-6.4.

My cost model's approach can also be used to decompose the cost effects into price vs. quantity, providing further insight on the role of each. Recall that using the decomposition in Section 5.1, cost equals quantity times price. Therefore, as long as expected quantity is positive under both networks, $dC_i=dQ_i\cdot dP_i$.[64] I can estimate regression (10) using quantity as the outcome variable to get an estimate of $d\hat{Q}_i=\exp\left(\hat{\gamma}_Q\left(Z_i\right)\right)$. The implied effect on prices is $d\hat{P}_i=d\hat{C}_i/d\hat{Q}_i=\exp\left(\hat{\gamma}_C\left(Z_i\right)-\hat{\gamma}_Q\left(Z_i\right)\right)$.

Table 12, columns (4)-(6) report estimates of this decomposition. Column (4) shows estimates for the subset of costs (inpatient and outpatient care) included in the decomposition; the estimates are quite similar to those for total costs. Interestingly, column (5) shows that most of the cost reductions represent a fall in *quantity* of care, with price reductions explaining a minority. While the average $\hat{\gamma}_C=-0.137$ (s.e. $=0.021$) corresponding to a 12.8% cost reduction, the average $\hat{\gamma}_Q=-0.105$ (s.e. $=0.020$) which is a 10% fall. Thus, quantity reductions account for about three-quarters of the fall in costs, while price reductions account for just one-quarter. The interactions with patient status also reveal interesting patterns. Both quantity and price reductions are largest for Partners patients (even controlling for other health measures), but quantity reductions still explain more than three quarters of the cost fall. For patients of other dropped hospitals, quantity falls but price increases, consistent with them substituting to higher-price providers.

[64] To see this use the notation of Section 2 to write $C_i(n)=Q_i(n)P_i(n)$ under network n (where 0 = narrower, 1 = broader). For a narrowing of the network, $dC_i=C_i(0)/C_i(1)=\left(Q_i(0)/Q_i(1)\right)\left(P_i(0)/P_i(1)\right)\equiv dQ_i dP_i$. Notice that this decomposition only works if *expected* quantity is positive under both networks (though *ex-post* realized quantity may be negative for some people), which is required for price to be well-defined. This seems like a reasonable assumptions for most people.

Table 12: Cost Model Estimates: Change in Cost with Narrower Network

Network Health x Post	Effect on Insurer Cost			Decomposition: Costs	Decomposition: Quantity	Decomposition: *Price*
	(1)	(2)	(3)	(4)	(5)	(6)
Average Effect	**-0.133*****	**-0.125*****	**-0.136*****	**-0.137*****	**-0.105*****	***-0.032***
	(0.018)	**(0.018)**	**(0.018)**	**(0.021)**	**(0.020)**	
Full Specification						
Constant	-0.133***	-0.089***	-0.236*	-0.367*	-0.387**	*0.020*
	(0.018)	(0.020)	(0.119)	(0.142)	(0.131)	
Patient of: Partners		-0.277***	-0.235***	-0.294***	-0.246***	*-0.048*
		(0.054)	(0.062)	(0.078)	(0.069)	
Other Dropped Hosp.		-0.071	-0.068	-0.076	-0.153*	*0.077*
		(0.070)	(0.070)	(0.073)	(0.074)	
Dist. to Partners: 0-2 miles *(omitted)*						
2-5 miles			0.004	0.029	0.030	*-0.001*
			(0.094)	(0.099)	(0.087)	
5-10 miles			0.072	0.124	0.099	*0.024*
			(0.098)	(0.102)	(0.090)	
10-20 miles			0.017	0.069	0.104	*-0.036*
			(0.097)	(0.106)	(0.094)	
20-30 miles			0.133	0.126	0.165	*-0.039*
			(0.099)	(0.104)	(0.092)	
>30 miles			0.032	0.095	0.148	*-0.053*
			(0.095)	(0.104)	(0.090)	
Other Interactions (summary)						
Age >= 45			0.084	0.214	0.233	*-0.019*
			(0.094)	(0.123)	(0.122)	
Risk score 40-80th%			-0.048	-0.039	-0.008	*-0.031*
			(0.074)	(0.094)	(0.089)	
Risk score >80th%			-0.039	-0.035	0.006	*-0.041*
			(0.067)	(0.086)	(0.080)	
Chronic illness			0.051	0.024	0.008	*0.016*
			(0.046)	(0.051)	(0.049)	
Cancer			-0.134**	-0.143**	-0.135*	*-0.008*
			(0.050)	(0.055)	(0.054)	
Number of Obs.		1,131,878			1,110,587	
Number of Individuals		128,496			125,572	

NOTE: The table reports estimates of cost changes due to Network Health's network narrowing in 2012, following the Poisson regression equation in (20). The estimates are of the $\gamma(Z_i)$ terms, which are approximately equal to percent effects on costs. More precisely, the multiplicative effect of the narrower network is $\exp(\gamma(Z_i))$, and the percent changes is $\exp(\gamma(Z_i)) - 1$. Columns (1)-(3) show estimates on total insurer costs for models with increasingly rich interactions. Column (4) shows the same specification as (3) but with a dependent variable of (inpatient/outpatient) costs covered by the price-quantity decomposition presented in Appendix C. Column (5) show estimates for changes in quantity, and (6) shows implied changes in prices.

D.5 Robustness Checks on WTP and Cost of Star Hospital Coverage

This appendix presents several modifications of ΔWTP and $\Delta Cost$ of Network Health's broader (2011)network that covers the star Partners hospitals in order to check the robustness of the finding in the body text (see Figure 8B) that ΔWTP is below $\Delta Cost$. See section 6.4 for the definition of these variables. Figure 12 replicates Figure 8B with several modified versions of these curves. In all cases, the baseline ΔWTP and $\Delta Cost$ curves are shown in green and black respectively with point markers. The modified curves are shown in curves without point markers. These modifications are:

1. Counting only quantity of care reductions in $\Delta Cost$ (Panel A): These define the cost reductions based only on reductions in quantity of care (price-standardized utilization), not reductions in price of care. Quantity reductions are estimated using the method in Appendix D.4 and specifically the estimates in column (5) of Table 12. Figure 12A shows both a high and low estimate of $\Delta Quantity$ calculated under different assumptions. The low estimate (gray curve) takes the estimates of proportional reductions in $\Delta Quantity$ and applies them to quantity of care *included in the cost decomposition* (inpatient and outpatient costs). This assumes that the one-third of costs not included in the decomposition do not change with the broader network, which likely generates a conservatively low estimate of $\Delta Quantity$.[65] Nonetheless, this low estimate of $\Delta Quantity$ is still substantially larger than ΔWTP by a factor of 2-3x. The high estimate (dark blue curve) assumes that the proportional reductions in $\Delta Quantity$ apply to total costs. This generates estimates quite similar to the baseline $\Delta Cost$ curve.

2. Recalculating $\Delta Cost$ using lower Partners prices (Panel B): This panel redefines cost reductions from the narrower network using lower Partners prices of care under the broader network. These lower prices could either reflect changes in hospital-insurer bargaining due to adverse selection or a lower social cost of care reflecting Partners' price markups.[66] To see how this works, note that $C_{ij}(1) = C_{ij}^{Partners}(1) + C_{ij}^{Other}(1)$, where the two terms reflect costs incurred at Partners and all other providers. Then $\Delta Cost_{ij} = C_{ij}(1) - C_{ij}(0)$. The modification recalculates $C_{ij}^{ALT}(1) = (1-\phi) \times C_{ij}^{Partners}(1) + C_{ij}^{Other}(1)$, where ϕ is a Partners price reduction factor. It then defines $\Delta Cost_{ij}^{ALT} = C_{ij}^{ALT}(1) - C_{ij}(0)$.[67] I consider price reductions (ϕ) of 10%, 25%, and 50%, reflecting a range of possible price reductions and/or markups.[68] Even with a 50% price reduction (an extreme

[65] The decomposition excludes items like prescription drugs, inpatient rehab, and some inpatient/outpatient costs that are paid in non-standard ways (see Appendix C). It seems likely that if included inpatient/outpatient costs fall substantially, these would also fall at least somewhat since their provision is also linked to the high-cost excluded Partners system. For instance, Partners owns a network of rehab hospitals (Spaulding Rehab), and costs may fall as patients substitute to other providers. Of course, it is also possible that non-included quantity of care moves in the opposite direction as included quantity (i.e., the two are substitutes), but this seems less likely. Against this possibility, the proportional reduction in total costs and included costs are quite similar – both are about 13-14% (see Table 12, columns (3) vs. (4) – which is consistent with the two moving in the same direction.

[66] The social value of these markups would depend on how the money is spent, which is an important but unclear issue. If used to increase hospital amenities (e.g., nicer buildings) or physician/administrator salaries, the social value might be less than dollar-for-dollar. If used to fund research and teaching, the social value might be more than dollar-for-dollar.

[67] This effectively assumes no change in Partners out-of-network prices under the narrower network that excludes it. This is conservative in that it will produce smaller estimates of $\Delta Cost^{ALT}$ than if I assumed $C_{ij}(0)$ also decreased.

[68] For context, a very rough calculation using state data on hospital costs per risk-adjusted discharge (CHIA, 2014b) suggests that the inpatient CommCare prices for MGH and Brigham & Women's (BWH) are marked up by about 20-30%

upper bound), the $\Delta Cost^{ALT}$ curve is still above WTP throughout the distribution.

3. Recalculating ΔWTP based on social marginal utility of money (Panel C): The final modification recalculates ΔWTP using a social marginal utility of money, which is a simple way to include a notion of equity in the welfare analysis. Note that baseline WTP is defined in equation (12) as the utility of the broader network (ΔV_i) divided by the marginal utility of money ($-\alpha(Z_i)$, the negative premium coefficient). We can define alternate $\Delta WTP^{ALT} = \Delta V_i / \left(-\alpha^{social}\right)$, where $-\alpha^{social}$ is a uniform social marginal utility of money (e.g., reflecting a cost of redistribution).[69] Panel C considers using two possible α^{social}: the average $\alpha(Z_i)$ among CommCare consumers and the 99th percentile $\alpha(Z_i)$ (i.e., close to the smallest in absolute value) which reflects the estimates for the highest-income (near 300% of poverty) and oldest (over age 60) consumers. The former does not affect ΔWTP much. The latter closes only a small part of the gap, with $\Delta Cost$ still 1.5-2.5x as large as ΔWTP^{ALT}. Note, however, that if I combine this high-end ΔWTP^{ALT} with the smallest version of $\Delta Cost^{ALT}$ with 50% lower Partners prices, the two are approximately equal. This illustrates the extreme modifications to WTP and costs that would be required to overturn the basic finding that WTP for the broader network falls short of costs.

relative to costs, while prices for the other five Partners hospitals are at or below costs. Outpatient care cost and markup data are not available, though the fact that Partners' outpatient care prices are not very high suggests they might be lower. Thus, 25% represents a high-end estimate of Partners' markup that assumes that the 20-30% inpatient markups for MGH and BWH apply to all care at Partners providers. Of course, hospital costs are known to be quite difficult to define and measure, so these figures should be taken to be very rough. Nonetheless, a 50% Partners price reduction is an extreme upper bound that would likely require Partners not just to cut markups (which are not "free," since markups are used to cross-subsidize other Partners activities) but also to make radical changes to how it delivers care.

[69] Note that a fully consistent cost-benefit analysis of the policy problem would need to explain why the government does not redistribute to CommCare enrollees (e.g., via lower premiums or cash checks) up to the point that their marginal utility of money equals the social cost of redistribution. This exercise is meant as illustrative, not a fully consistent policy analysis of equity and redistribution.

Figure 12: Robustness Analysis: ΔWTP and ΔCost of Broader Network

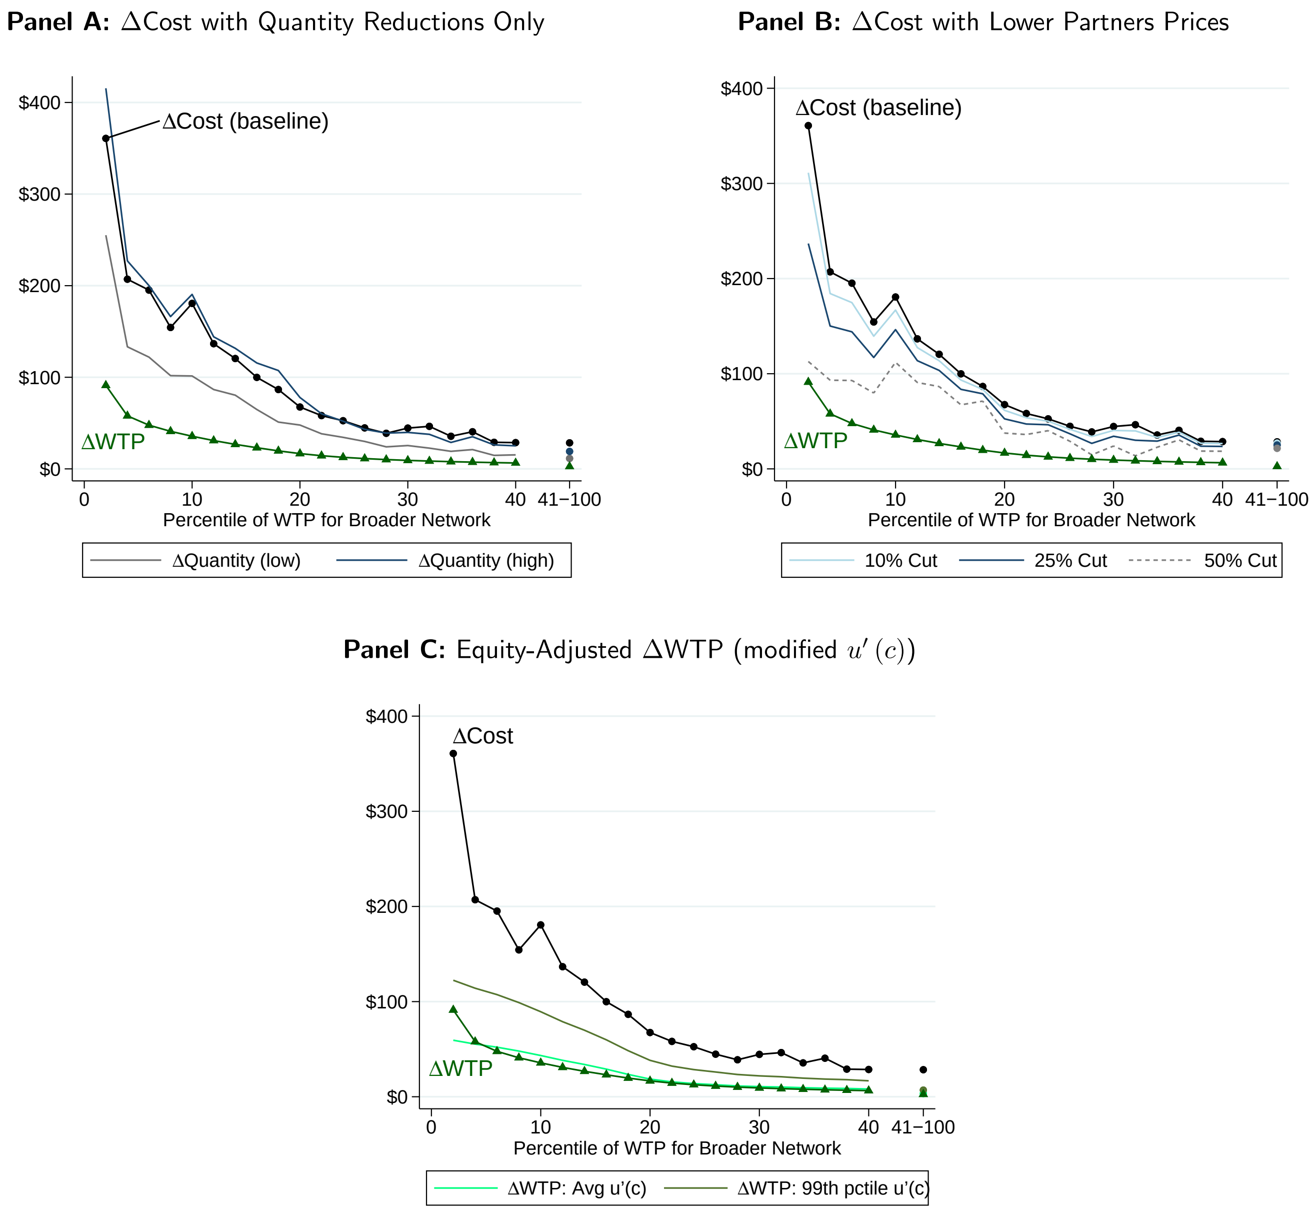

NOTE: These figures replicate Figure 8 in the body text with various modifications to $\Delta Cost$ and ΔWTP. See the note to Figure 8 and the text of Appendix D.5 for additional information describing the definition of these curves. The overall conclusion is that $\Delta Cost$ still exceeds ΔWTP under a variety of alternate definitions of the two.

Appendix E: Robustness and Additional Figures/Tables

E.1 Cost Breakdown by Group for Network Health Enrollees

Table 13: Analysis of Costs for Network Health Enrollees in 2011 (Stayers and Switchers)

Enrollee Group	All Network Health Enrollees in 2011 (Switchers + Stayers)			Switching Out Choices		Risk Adj. Cost Among Switchers Out	
	Raw Cost	Risk Adj. Cost	Share of Enrollees	Switching Rate	Share of Switchers	2011	2012
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
All Enrollee Groups	$366	$346	100%	11%	100%	$508	$452
By Prior-Year Care							
Partners Hospitals	$701	$564	18%	45%	67%	$572	$475
Other Dropped Hospitals	$487	$386	8%	24%	17%	$375	$372
All Other Enrollees	$273	$274	74%	3%	16%	$333	$422
By Distance to Partners Hospital							
0-5 miles	$383	$363	23%	22%	46%	$469	$478
5-25 miles	$371	$354	36%	12%	36%	$512	$399
> 25 miles	$353	$329	41%	5%	18%	$583	$497

NOTE: The table shows statistics about continuing enrollees in Network Health in 2011, including both individuals who stick with the plan in 2012 ("stayers") and those who switch to another plan in 2012 ("switchers out") when the network changes. The top row (highlighted in gray) shows overall average statistics, and the following panels show subgroup averages by prior-year outpatient care use and by enrollee distance to the nearest Partners hospital. Columns (1)-(3) show statistics (raw cost, risk adjusted costs, and the share each group represents) for all switchers and stayers together. Columns (4)-(5) show switching rates and shares of switchers each subgroup represents. Columns (6)-(7) show average risk-adjusted costs for 2011 and 2012 conditional on switching out.

E.2 Robustness Analyses on Adverse Selection Evidence

Figure 13: Changing Risk Selection for Network Health among New Enrollees

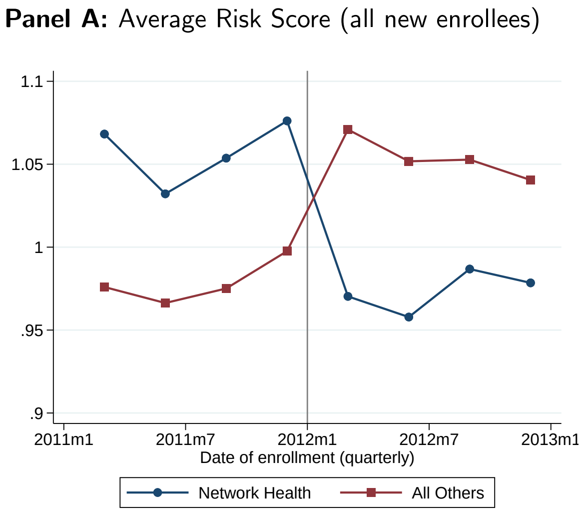

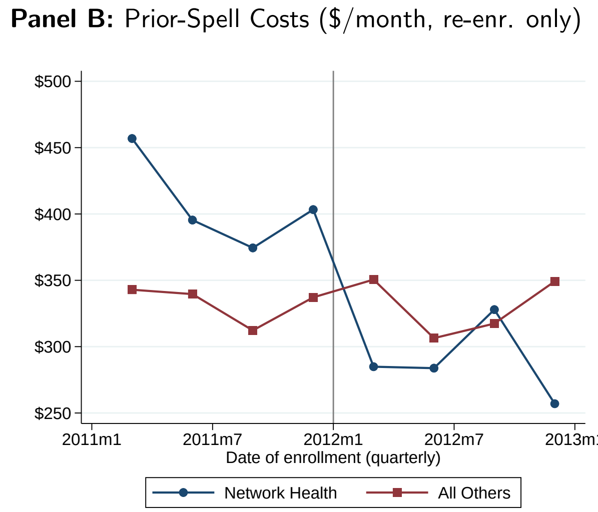

NOTE: These figures show evidence that shifts in new enrollee demand for Network Health at its 2012 network narrowing were correlated with proxies for cost in a way suggesting more favorable selection. Each point on the figures shows an average value for above-poverty new enrollees joining in a given bimonthly period who select Network Health (blue series) and all other plans (red series). The sample is restricted to above-poverty enrollees who are not subject to the 2012+ limited choice policy. Panel A shows average CommCare risk score (for all new enrollees). The average risk of Network Health's enrollees fell at the start of 2012 while that of other plans rose, suggesting a shift of high-risk enrollees from Network Health to other plans. Panel B shows prior-spell average costs (in $ per month) with the sample limited to re-enrollees who have a prior CommCare enrollment spell. The average cost of Network Health's enrollees falls at the start of 2012, while that of other plans is relatively flat.

Figure 14: Plan Switching and Selection for CeltiCare (Drops Partners in 2014)

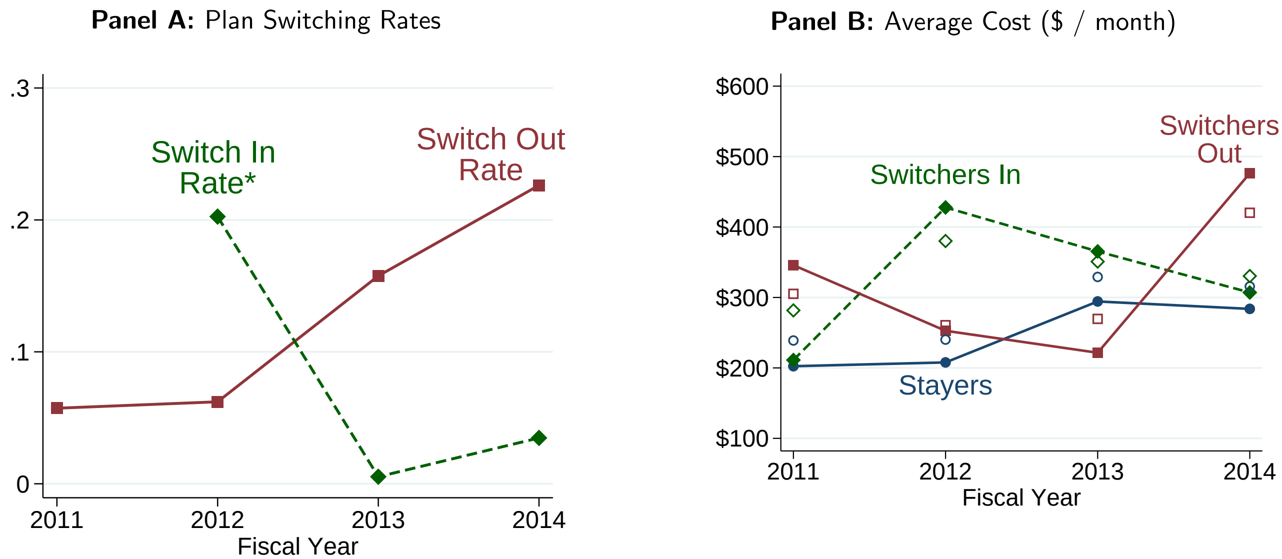

* Panel A excludes the 2011 switching in rate for CeltiCare to avoid blowing up the y-scale.

NOTE: These figures show switching rates for CeltiCare (Panel A) and average prior-year costs for CeltiCare enrollees (Panel B, in $ per month) in each year's open enrollment. CeltiCare drops the Partners Healthcare system from its network in 2014. These plots are analogous to Figure 2 in the main text, which show switching and selection for Network Health. See the note to Figure 2 for additional description. The current figure shows that similar adverse selection patterns occur for CeltiCare when it excludes Partners from network.

Figure 15: Switching Out Rates for CeltiCare (Drops Partners in 2014), by Enrollee Characteristics

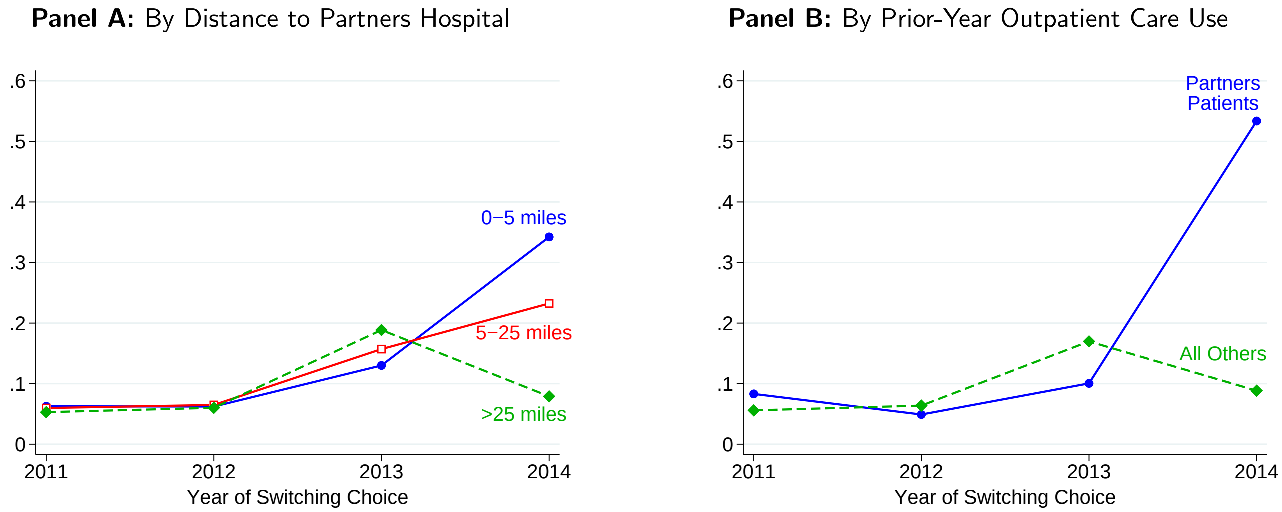

NOTE: These figures show switching out rates for CeltiCare enrollees by variables likely to correlate with demand for Partners, which is dropped from the plan's network in 2014. Panel A shows switching rates by enrollee distance to the nearest Partners hospital; Panel B shows switching rates by prior-year use of Partners for (non-emergency room) outpatient care. These plots are analogous to Figure 3 in the main text, which show switching for Network Health. See the note to Figure 2 for additional description.

Figure 16: CeltiCare's Market Share among New Enrollees (Drops Partners in 2014)

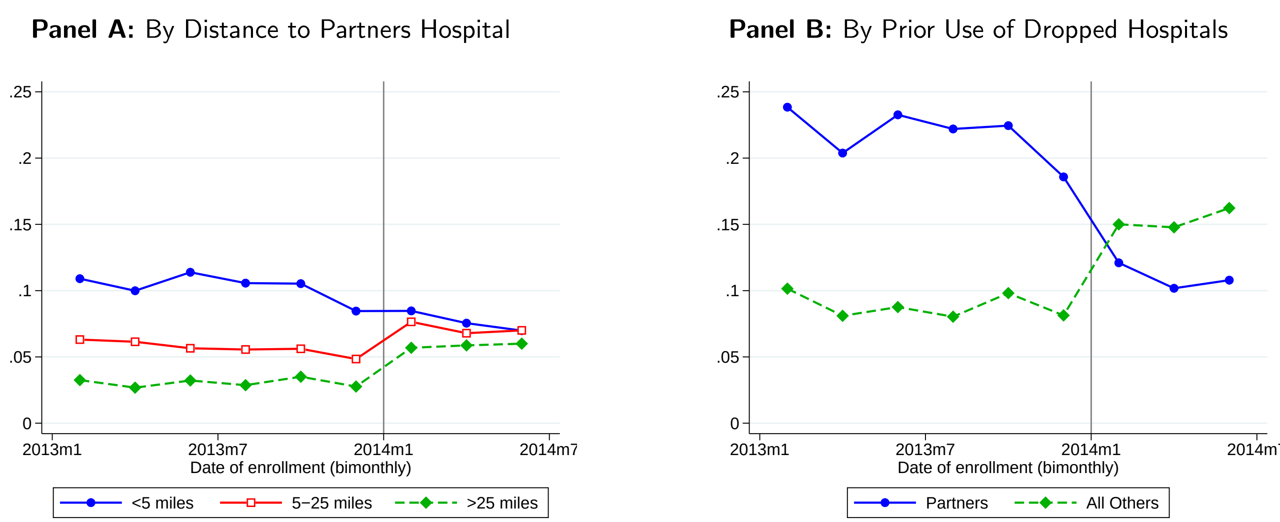

NOTE: These figures show evidence of changes in new enrollees' demand for CeltiCare in 2014 (when it drops Partners from network) that are correlated with valuation for Partners providers. (The plots are analogous to Figure 4 in the main text, which studies Network Health's network change in 2012.) Each point on the figures is the share who choose CeltiCare among above-poverty new enrollees joining the exchange in a given (bimonthly) period. The sample is restricted to above-poverty enrollees who are not subject to the 2012+ limited choice policy. Panel A divides enrollees by proximity to the nearest Partners hospital. Panel B divides enrollees by use of the Partners hospitals during a prior enrollment spell (with the sample limited to re-enrollees who have a prior spell). The slightly "early" decline in the market share for Partners patients (in the final period of 2013) reflects the fact that the network change was announced prior to its enactment at the start of fiscal year 2014.

Figure 17: Plan Switching and Selection for Network Health: Zero-Premium Enrollees

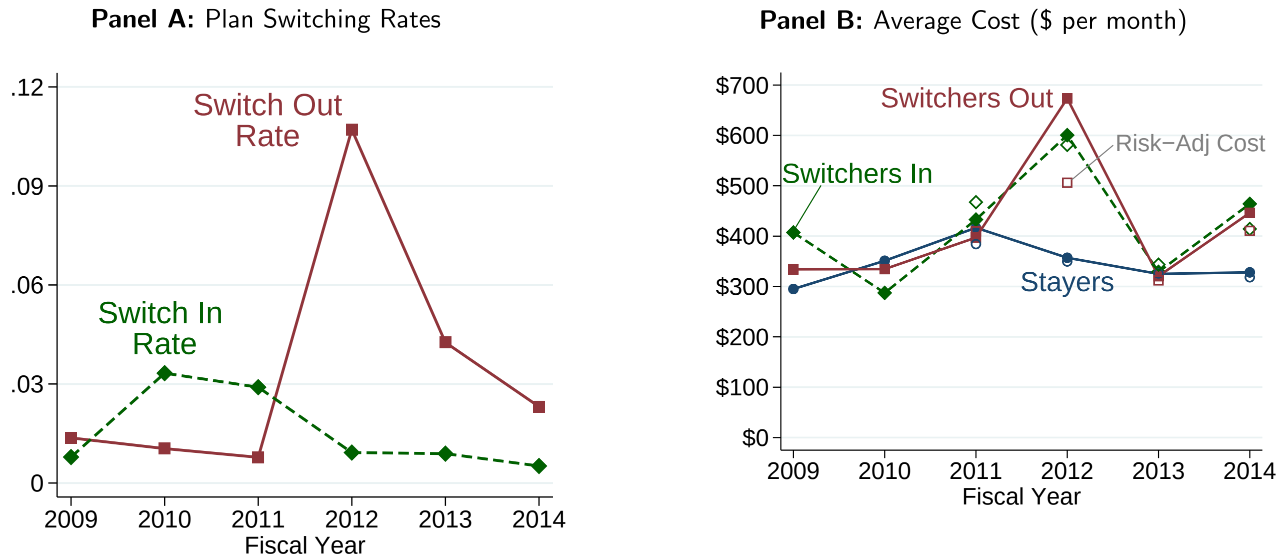

NOTE: These figures show switching and selection patterns for *zero-premium* (below-poverty) Network Health over time and especially around its 2012 network narrowing. The graphs are exactly analogous to Figure 2 in the main text but with the sample limited to below-poverty enrollees who do not pay premiums. See the caption to Figure 2 for additional information.

Figure 18: Plan Switching Out Rates for Network Health: Zero-Premium Enrollees

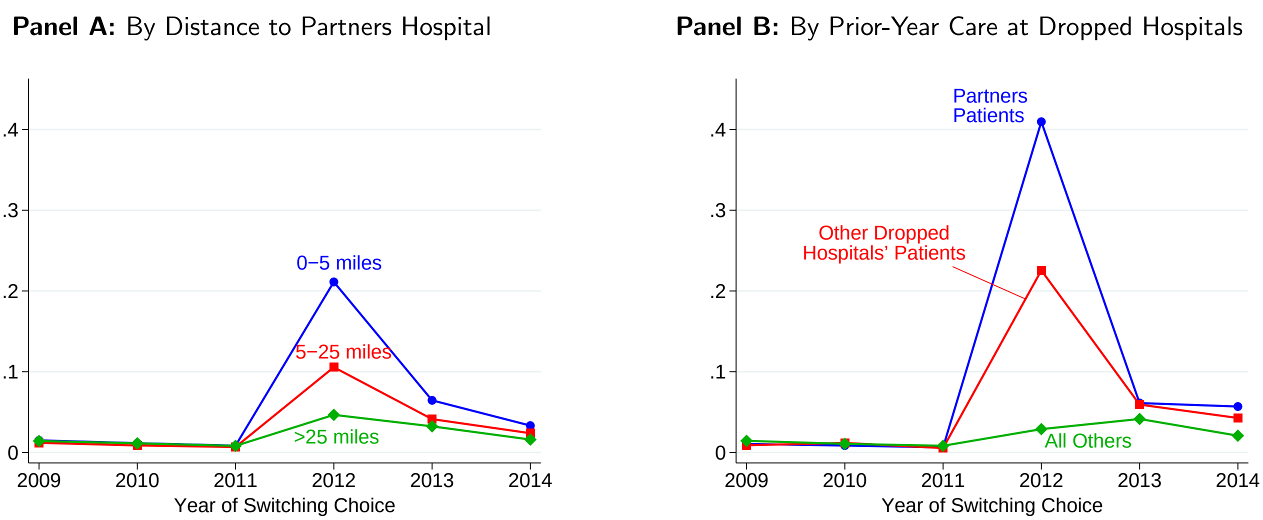

NOTE: These figures show switching out patterns for *zero-premium* Network Health enrollees around its 2012 dropping of Partners and several other hospitals. They are exactly analogous to Figure 3 in the main text but with the sample limited to below-poverty enrollees who do not pay premiums. See the caption for Figure 3 for additional information.

Figure 19: Monthly Rate of Exiting the Exchange, Network Health Enrollees

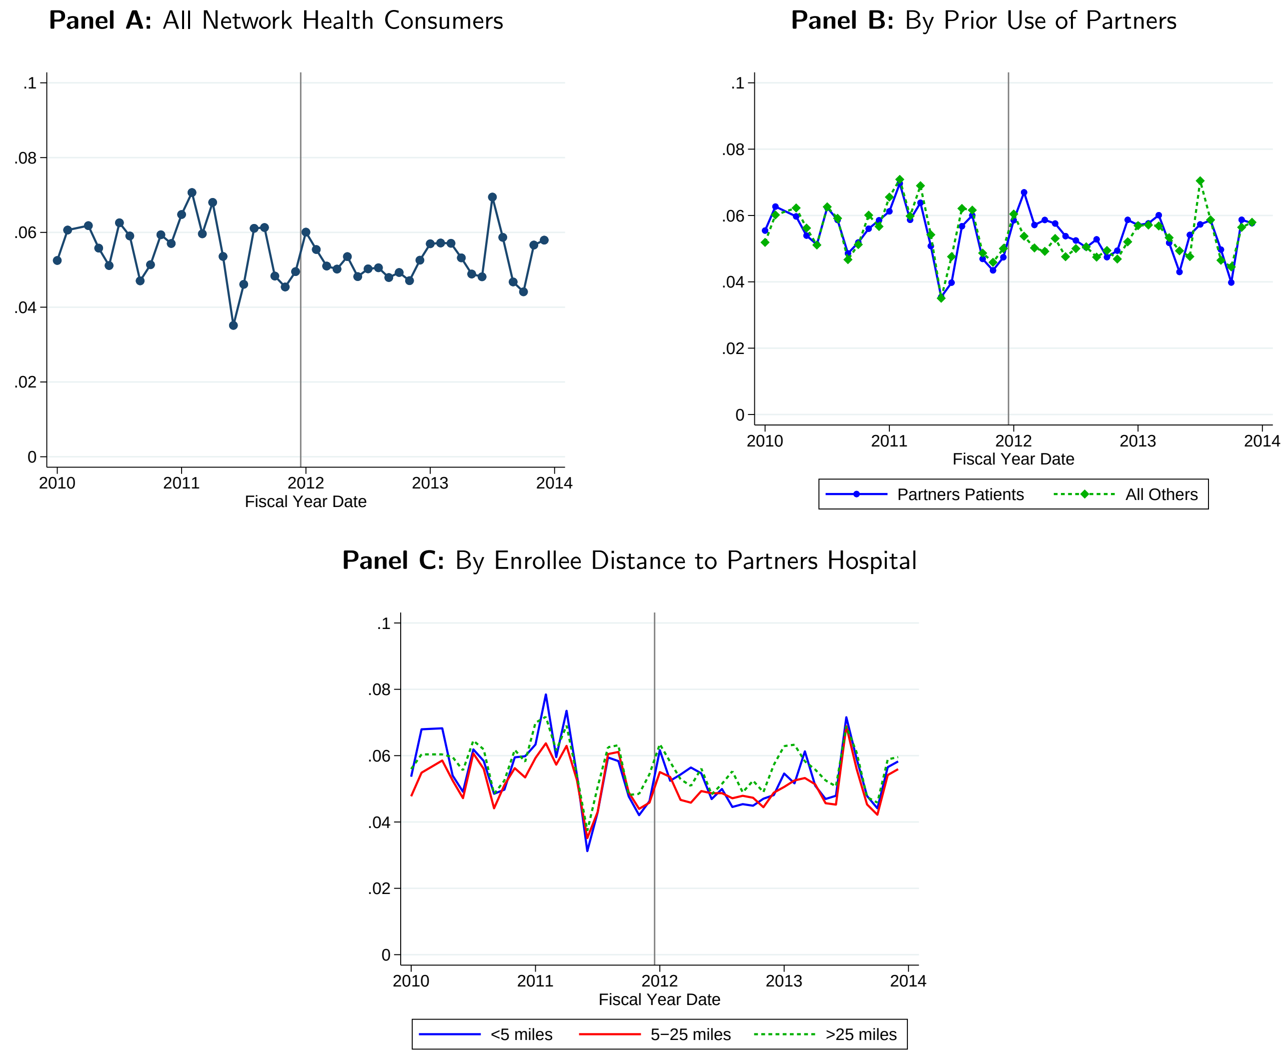

NOTE: The figure provides evidence on a key assumption in the plan choice model: that Network Health's network narrowing in 2012 does not affect whether consumers participate in the exchange (no "extensive margin" response). The figure plots the share of Network Health's existing enrollees who exit the exchange in each month from 2010-2013. If the network narrowing in 2012 led to an extensive margin response, we would expect to see a jump upward in the exit rate at the start of 2012. There is little evidence of this either for Network Health enrollees overall (panel A) or when broken down by factors that strongly predicted plan switching: Partners patients vs. others (panel B) or enrollee distance to a Partners hospital (panel C).

Table 14: Switching Rate Out of Network Health, by Recency of Last Provider Visit (Estimates)

Recency of Latest Visit to Provider	Probability Switch Out of Network Health in 2012					
	Raw Probabilities (no controls)		Medical Risk Controls		Medical Risk and Distance Controls	
	(1)		(2)		(3)	
	Prob.	*(S.E.)*	*Prob.*	*(S.E.)*	*Prob.*	*(S.E.)*
Partners Patients						
1 month	0.730	(0.016)	0.726	(0.016)	0.713	(0.017)
2 months	0.651	(0.020)	0.641	(0.021)	0.624	(0.022)
3-4 months	0.545	(0.019)	0.541	(0.020)	0.522	(0.021)
5-6 months	0.565	(0.027)	0.572	(0.028)	0.562	(0.029)
7-9 months	0.448	(0.026)	0.455	(0.027)	0.447	(0.028)
10-12 months	0.433	(0.033)	0.449	(0.034)	0.432	(0.034)
13-18 months	0.339	(0.026)	0.349	(0.027)	0.364	(0.028)
19-24 months	0.226	(0.022)	0.229	(0.023)	0.246	(0.025)
>24 months	0.186	(0.016)	0.180	(0.016)	0.194	(0.017)
Other Dropped Providers' Patients						
1 month	0.469	(0.036)	0.446	(0.038)	0.405	(0.039)
2 months	0.375	(0.043)	0.367	(0.045)	0.357	(0.047)
3-4 months	0.292	(0.031)	0.304	(0.034)	0.279	(0.034)
5-6 months	0.290	(0.055)	0.273	(0.056)	0.240	(0.054)
7-9 months	0.183	(0.038)	0.179	(0.039)	0.171	(0.039)
10-12 months	0.137	(0.040)	0.112	(0.036)	0.120	(0.039)
13-18 months	0.123	(0.031)	0.112	(0.030)	0.105	(0.029)
19-24 months	0.071	(0.028)	0.066	(0.027)	0.062	(0.026)
>24 months	0.213	(0.031)	0.192	(0.031)	0.165	(0.029)
All Other Patients						
1 month	0.084	(0.003)	0.076	(0.003)	0.063	(0.003)
2 months	0.081	(0.004)	0.076	(0.004)	0.062	(0.003)
3-4 months	0.075	(0.004)	0.072	(0.004)	0.059	(0.003)
5-6 months	0.055	(0.005)	0.054	(0.005)	0.044	(0.004)
7-9 months	0.047	(0.004)	0.046	(0.004)	0.036	(0.004)
10-12 months	0.060	(0.006)	0.061	(0.007)	0.049	(0.005)
13-18 months	0.038	(0.005)	0.040	(0.006)	0.030	(0.004)
19-24 months	0.041	(0.007)	0.044	(0.008)	0.032	(0.006)
>24 months	0.031	(0.005)	0.033	(0.005)	0.024	(0.004)

NOTE: The table reports estimates corresponding to Figure 5 in the text. Individuals are categorized into Partners patients (top panel), patients of another dropped hospital (middle panel), and all other patients (bottom panel) based on prior physician office visits in the claims data. Individuals with no prior office visits (about 19%) in the data are excluded. Among the remaining sample, 13% have a prior Partners visit and 4% have a prior visit to another dropped hospital's physician, with a small number of overlaps (0.3%) classified as Partners patients. The table shows rates of switching out of Network Health in 2012 by recency (as of the start of 2012) of the latest physician office visit to the indicated provider (e.g., Partners for the Partners patients). The numbers shown are predicted probabilities for bins of recency (using Stata's "margins" command) from logit regressions with various controls, evaluated at control variable means. Column (1) has no control variables; column (2) controls for demographics (age, gender, income group) and medical risk variables (chronic condition dummies and ventiles of HCC risk score); column (3) additionally controls for distance to Partners and other dropped providers, using the distance categories in Table 2. Separate regressions are run for each patient group.

E.3 Additional Estimates of Causal Cost Effects (Moral Hazard)

Figure 20: Event Study: Cost Reductions after 2012 Network Change

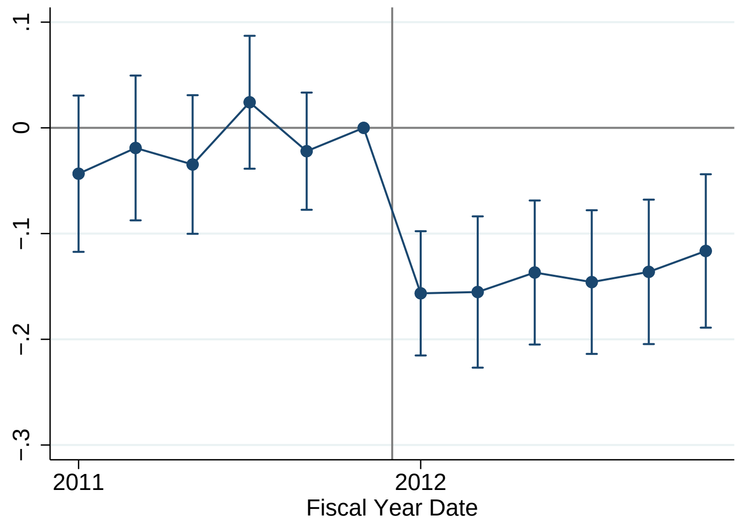

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (10) in Section 5.3 (and corresponding to Panel A of Figure 6). Estimates are from a Poisson regression with individual fixed effects and capture the cost differences between stayers in Network Health from 2011-12 versus stayers in other plans (control group), relative to the omitted period (the final bimonthly period of 2011). Poisson coefficients are roughly interpretable as percent differences; more precisely the percent difference is $\exp(\gamma) - 1$. The figure confirms the presence of parallel pre-trends and a sharp and persistent fall in costs of about 10-15% during 2012.

Figure 21: Event Study: Cost Reductions after 2012 Network Change, by Partners Patients

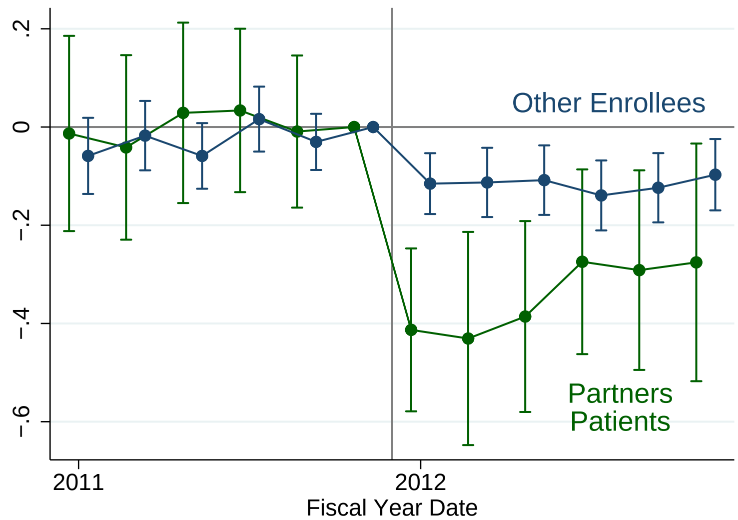

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (10) in Section 5.3, with separate interactions for γ with Partners patients (green series) versus other enrollees (blue), corresponding to Panel B of Figure 6. See note to Figure 20 for additional information on the setup and interpretation of coefficients. This figure corresponds the presence of parallel pre-trends for both groups and a share cost reduction in 2012 that is much larger for Partners patients.

Figure 22: Cost Reductions after 2012 Network Change, by Distance to Partners Hospital

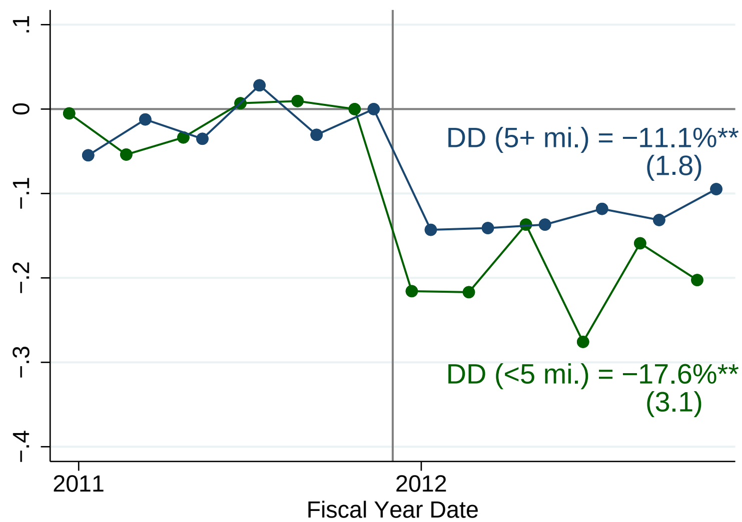

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (10) in Section 5.3, with separate interactions for γ with people living within 5 miles of a Partners hospital (green series) versus those living 5+ miles away (blue). See note to Figure 20 for additional information on the setup and interpretation of coefficients. Confidence intervals are suppressed because they are sufficiently wide to make it difficult to see the two series. The overall DD coefficients and their confidence intervals are reported. These are suggestive of a larger cost reduction for people living within 5 miles of a Partners hospital, but note that the difference is not statistically significant.