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

Mark Shepard*

August 25, 2021

Abstract

Health insurers increasingly compete on their networks of medical providers. Using data from Massachusetts' insurance exchange, I find substantial adverse selection against plans covering the most prestigious and expensive "star" hospitals. I highlight a theoretically distinct selection channel: consumers loyal to star hospitals incur high spending, conditional on their medical state, *because* they use these hospitals' expensive care. This implies heterogeneity in consumers' incremental costs of gaining access to star hospitals, posing a challenge for standard selection policies. Along with selection on unobserved sickness, I find this creates strong incentives to exclude star hospitals, even with risk adjustment in place.

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 Olivia Zhao, 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 Marcella Alsan, 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 limit coverage to a subset with whom they have negotiated contracts. "Narrow network" plans have proliferated in market-based public programs like the Affordable Care Act (ACA), Medicaid managed care, and Medicare Advantage that let enrollees choose among competing plans. 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.[1] Within the selection literature, there is no direct evidence on the connection between networks and selection incentives.[2] 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, Forthcoming).

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

To provide evidence, I study Massachusetts' pre-ACA health insurance exchange, a model market for the ACA. Using variation in coverage of the state's top star hospital system, I find substantial selection incentives to exclude the star providers. These incentives persist despite sophisticated risk adjustment intended to offset adverse selection. Investigating the mechanisms, I find a key role for both unobserved medical risk and a non-standard channel: people who demand star hospital coverage have higher costs *because* they use its expensive care. This channel creates selection on moral hazard and poses challenge for risk adjustment and standard policy responses to selection.

The paper has two main contributions. The first is the basic finding of adverse selection on star hospital coverage. The Massachusetts exchange setting is ideally suited to this topic because plan financial benefits (cost sharing and covered services) are standardized, letting me study plans that are

[1] See Einav, Finkelstein and Levin (2010a) 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.

[2] 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.

nearly identical except for networks. Moreover, there is variation in coverage of the state's top star hospital system: Partners Healthcare, which is both the state's largest health system and includes two nationally top-ranked hospitals (Mass. General and Brigham & Women's hospitals).

The main evidence comes from a large plan that drops Partners as part of shifting towards a narrow-network, low-price strategy. I use this as a natural experiment to test for selection. Just after the exclusion, the plan sees a large exodus of high-cost consumers who live near a Partners hospital and/or who are existing patients of a Partners provider. About 45% of Partners patients switch out of the plan – a more than six-fold increase, and strikingly high given well-known consumer inertia (Handel, 2013; Ericson, 2014). Relative to stayers, switchers had 108% higher costs, and 60% higher after risk adjustment, levels that made them unprofitable. Meanwhile, the plan also benefited from an influx of low-cost consumers attracted by the plan's lower price. These patterns illustrate the competitive logic of adverse selection. Dropping the star hospitals led many people to leave the plan, but this *improved* its bottom line (while raising rivals' costs) because the switchers were high-cost and unprofitable.

My paper's second contribution is to analyze the mechanisms underlying adverse selection on star hospital coverage. My main conceptual point is that consumers incur high spending for two reasons, or along *two cost dimensions*. The standard dimension is *medical risk* (or sickness). Medical risk reflects patient attributes that predict greater illness risk and use of care, regardless of the provider. Most analyses of adverse selection implicitly assume sickness is the main or only reason for cost variation.

But when plans compete on networks, a second cost dimension is also relevant: variation due to *use of expensive providers*. This dimension arises from the interaction of two forms of heterogeneity. First, providers vary in their overall "expensiveness." Spending for a given illness is not mechanical but is influenced by the provider, both through treatment decisions (quantity of care) and through prices per service. Both treatment intensity and prices vary widely (Cooper et al., 2019) and tend to be high at star hospitals (Newhouse, 2003; Ho, 2009). Although expensiveness is a provider attribute, it interacts with a second form of heterogeneity: varying consumer demand for providers. Demand varies for many reasons, including medical considerations but also (non-medical) preferences. Putting these two together, patients with higher demand for expensive providers will be differentially costly to insurers, even conditional on medical risk.

I formally define and analyze the properties of selection along these two cost dimensions in Section 2. In some ways they are similar. Both imply higher insurer average costs and may discourage coverage of an expensive hospital and/or push expensive hospitals to accept lower prices.[3]

However, in other ways selection on expensive provider use is different. The key economic difference is how it interacts with the network. Whereas medical risk is a (largely exogenous) patient attribute, using expensive providers is endogenous and can be avoided with a narrower network that steers patients to cheaper providers. Access to expensive star hospitals can be thought of as an "extra" benefit (on top of the minimum required network), which benefits patients but also raises costs. The

[3]Although I do not model bargaining, I argue that adverse selection reduces star hospital leverage and can put downward pressure on their prices – a point also noted by Ho and Lee (2017; 2019) and related to their discussion of the "recapture effect." This could lead to a more desirable outcome – which I do not see in my empirical setting – of lower provider prices without network exclusion.

provider cost channel shows up in consumers' *incremental costs* of access to a broader network – or the "moral hazard" response to the network. The selection challenge is that incremental costs vary widely across consumers based on their demand for the star hospital. For instance, costs may increase negligibly for someone living hundreds of miles away from the star hospital (low demand) but increase substantially for someone living next door to it (high demand). This sets up the conditions for "selection on moral hazard" (Einav et al., 2013), a key prediction of this second cost channel.

This core economic difference implies several others. First, the expensive providers channel is especially likely to create adverse selection because the same provider demand driving high costs (via star hospital use) also affects plan choice. This creates the link between demand and costs that is the hallmark of adverse selection. Second, even excellent risk adjustment is unlikely to offset selection on this channel because of the role of preferences in star hospital demand. Finally, the connection between selection and moral hazard (and thus, *selection on moral hazard*) makes policy responses challenging. Instead of being a technical problem to be "fixed" with subsidies or mandates, adverse selection is tied up in the difficult tradeoff between generous coverage and moral hazard (Einav et al., 2016).

I use the Massachusetts data to gain insight on the role of these two cost dimensions for adverse selection on coverage of the star Partners hospitals. I start by analyzing the determinants of demand for Partners. A natural question is whether to think of Partners as a vertically superior provider (as its *U.S. News* rankings suggest) or as a horizontally differentiated provider that happens to be expensive. The vertical model suggests demand that is concentrated among the sick, while the horizontal model suggests a larger role for preferences. In practice, I find evidence for *both* of these stories. Demand for Partners is strongly correlated with being sicker (e.g., being in the top 5% of risk scores) and with preference measures (e.g., distance to Partners). But quantitatively, preferences appear to explain more of the variation. Distance, which is just one determinant of preferences, accounts for 56-69% of the explained variation in demand measures, versus 2-8% explained by "observed risk" (variables used in risk adjustment) and another 28-35% by a richer set of measures derived from claims data ("unobserved risk"). Moreover, there appears to be a large role for unobserved preferences and/or provider loyalty, as suggested by the strong power of patient-doctor relationships in explaining demand.

This mixture of preferences and sickness driving star hospital demand suggests a policy dilemma. If demand were purely about sickness, regulators might want to subsidize or mandate star hospital coverage, even at extra cost, just as they mandate other "essential health benefits" used by the sick. If demand were purely preferences, they might be comfortable letting coverage unravel. The mixture of preferences and sickness, instead, suggests a difficult tradeoff.

I next use the claims data and the 2012 network change to disentangle the sources of high costs among people who value the star hospitals – i.e., the sources of adverse selection. I find a role for both the standard medical risk and non-standard expensive providers dimensions. High utilization linked to medical risk explains just over half (53%) of switchers' higher costs, with most of this being "unobserved risk" not captured by the exchange's risk adjustment. Use of high-price Partners providers explains a meaningful share (22%) of inpatient costs, while residual quantity (not explained by risk) is more important for outpatient costs. A key piece of evidence for the role of the expensive providers channel comes from estimating within-person cost changes for stayers who remain in the plan that drops Partners between 2011 and 2012. I find a sharp 15% cost reduction for stayers, occurring through both lower prices and quantity of care. Cost reductions are much larger for Partners patients (about 30%, or \$175 per month) than for other enrollees (9%, or \$30 per month), consistent with the key prediction of heterogeneity in incremental costs.

If risk adjustment breaks down, should regulators subsidize or mandate coverage of star hospitals? My analysis highlights the difficult tradeoffs involved with these policies. On the one hand, demand for star hospitals partly reflects sickness. Therefore, promoting broader networks differentially helps the sick, whose access regulators may want to ensure. On the other hand, star hospital coverage involves higher costs. Most of the adverse selection I find is driven by selection on incremental costs. Indeed, my model estimates suggest that incremental costs for Partners coverage rise even more steeply than incremental willingness to pay (WTP), creating the conditions for inefficient or even "backward" sorting highlighted in recent work (Bundorf et al., 2012; Marone and Sabety, 2021).[4] Consumer WTP for Partners falls short of incremental costs for the entire distribution, suggesting that excluding Partners was efficient given its observed cost structure.[5]

This paper's results are important for several reasons. First, they show the continued relevance of adverse selection, even in markets that try to address it through regulation and risk adjustment. They suggest a general mechanism – preferences for using expensive providers – through which selection can persist. Second, they illustrate the powerful economic forces pushing towards narrower networks in individual health insurance markets like the ACA exchanges. Finally, they show the challenge when selection and moral hazard are linked. Selection on moral hazard is not just a technical problem to be "fixed" with smarter risk adjustment or subsidies; rather it is an economic problem tied into fundamental tradeoffs between costs, quality, and access to top providers.

The paper proceeds as follows. Section 2 presents a model formalizing the paper's main ideas. Section 3 introduces the setting and data. Sections 4-5 show reduced form evidence and analyze the mechanisms for costs. Section 6 presents and analyzes a structural model, and Section 7 concludes.

2 Conceptual Model

I start with a model to formalize the mechanisms for adverse selection on provider networks. The model highlights two dimensions by which consumers may have high costs: (1) medical risk and (2) costs due to use of expensive providers for care.

2.1 Model Setup and the Selection Incentive

Consider an insurance market where single-plan insurers compete on premiums and provider networks. Based on the empirical setting, I focus on the decision of a single insurer j to cover vs. exclude a top

[4] Given the role of incremental costs and high prices, the most natural policy responses target use of expensive providers. These might include physician incentives to consider costs when making referrals (Song et al., 2012; Ho and Pakes, 2014) or higher "tiered" copays for expensive providers (Prager, 2020). Of course, the latter would need to be carefully weighed against losses in risk protection, especially for a low-income population.

[5] Although part of these incremental costs reflect the star hospitals' high price markups (which are a transfer, not a real cost), I find that WTP is still well below "adjusted" incremental cost curves that apply reductions to Partners prices of up to 50%.

star hospital, h^S, at a fixed set of hospital prices.[6] The star hospital is assumed to be highly valued by many consumers but also expensive. Aside from coverage of h^S, I assume the rest of insurer j's network, and the networks and negotiated hospital prices of all other insurers, are held fixed. However, both j and other plans can observe the network decision and adjust premiums in response. Consumers then follow by choosing among available plans, and when sick, choosing providers and incurring costs.

Let $n_j \in \{0,1\}$ denote whether plan j chooses to cover the star hospital, and $P(n_j)$ be the premiums that follow under each n_j choice. Let $D_{ij}(n_j)$ indicate whether consumer i chooses plan j, given its network decision $n_j \in \{0,1\}$ and the resulting premiums, $P(n_j)$. A key outcome for selection is consumers' *change in demand* in response to the network shift, or $\Delta D_{ij} \equiv D_{ij}(1) - D_{ij}(0)$. It is natural to expect that demand changes will align with consumers' value for access to the star hospital.[7]

Likewise, let $C_{ij}(0)$ and $C_{ij}(1)$ be expected insurer j costs for consumer i under the narrower and broader network. I call $\Delta C_{ij} \equiv C_{ij}(1) - C_{ij}(0)$ the "incremental cost" on consumer i of the broader network. Because h^S is expensive, we expect $\Delta C_{ij} \geq 0$. Importantly, ΔC_{ij} is likely to vary widely across consumers and may be correlated with star hospital demand.

The exchange seeks to mitigate adverse selection through risk adjustment. Although the plan must charge a single premium P_j for all consumers, the regulator adjusts revenues so the plan receives $\varphi_i P_j$ for consumer i, where $\varphi_i \equiv E(C_{ij}|Z_i)/\overline{C}$ is a "risk score" that estimates i's relative costliness based on medical observables Z_i. Therefore, the profitability of i under network n_j equals $\varphi_i P_j - C_{ij}(n_j)$. Following Curto et al. (2021), it is useful to factor out φ_i and write total profits as:

$$\pi_j(n_j) = \sum_i \left[P_j(n_j) - C_{ij}^{RA}(n_j)\right] \cdot \varphi_i D_{ij}(n_j) \tag{1}$$

where $\varphi_i D_{ij}(n_j)$ is risk-scaled demand and $C_{ij}^{RA}(n_j) \equiv C_{ij}(n_j)/\varphi_i$ is risk-adjusted costs. The outcome of interest is how j's *profits change* when it covers the star hospital, which can be decomposed as:

$$\Delta\pi_j = \underbrace{\sum_i \left[\Delta P_j - \Delta C_{ij}^{RA}\right] \cdot \varphi_i D_{ij}(0)}_{\textbf{(1) Fixed Enrollment: } \text{Premium and Cost change}} + \underbrace{\sum_i \left[P_j(1) - C_{ij}^{RA}(1)\right] \cdot \varphi_i \Delta D_{ij}}_{\textbf{(2) Selection: } \text{Profitability of marginal enrollees}} \tag{2}$$

Term (1) represents the impact of the plan's premium change and incremental costs (moral hazard) due to the broader network, holding enrollment fixed. Term (2) represents the selection incentive, which equals the profitability of marginal enrollees who select into/out of the plan due to the network/premium changes (i.e., $\Delta D_{ij} \neq 0$). There is an *adverse selection incentive* if people who select in ($\Delta D_{ij} > 0$) have high risk-adjusted costs and/or people who select out ($\Delta D_{ij} < 0$) have low risk-adjusted costs, where high/low are relative to $P_j(1)$.

[6] A broader model would have several stages: hospital-insurer network and price bargaining, followed by premium setting, then consumer plan choice, and then consumer hospital choice (e.g., Ho and Lee 2017; 2019). My setup focuses on a small part of the bargaining game to highlight the role of adverse selection.

[7] This is easiest to see in the (likely) case that j raises its premium when it covers the star hospital. Then consumers who highly value star hospital access will be more eager to shift toward j ($\Delta D_{ij} > 0$) despite the higher fee, while consumers with lower values for it will be more likely to shift away ($\Delta D_{ij} < 0$). More generally, this follows naturally in any choice model where consumers have heterogeneous preferences over star hospital coverage and other differentiated plan attributes.

2.2 Two Dimensions of Costs and the Limits of Risk Adjustment

Why would there be adverse selection incentives, given the regulator's attempts to offset it with risk adjustment? A key reason is that cost variation arises not just from medical risk but also from *varying demand for* (and use of) the expensive star hospital. To understand the logic, consider first a simpler "risk-only" model in which risk adjustment *does* work well. Suppose that consumers face risks, r_{id}, of various illnesses $d \in \{1, ..., D\}$, with illness d resulting in expected costs ω_d. Define $R_i \equiv \sum_d r_{id}\omega_d$ as overall risk for consumer i. Additionally, let κ_j be a constant factor capturing insurer j's cost structure; for instance, this might capture differences in plan actuarial value or administrative efficiency. In the risk-only model, risk-adjusted costs equal:

$$\textbf{Risk-Only model:} \qquad C_{ij}^{RA} = \underbrace{(R_i/\varphi_i)}_{\text{Unobserved risk}} \times \underbrace{\kappa_j}_{\text{Plan effect (constant)}} \tag{3}$$

In this model, risk-adjusted costs vary across consumers only if there is unobserved risk – that is, if risk scores (φ_i) do not fully capture true risk (R_i). The goal of regulators is primarily *statistical*: improving measurement and modeling so risk scores get closer to perfectly capturing risk (i.e., $R_i/\varphi_i \to 1$). If this occurs, $C_{ij}^{RA} = \kappa_j$. Differences in cost structure pass through into risk-adjusted costs – preserving insurers' incentives to improve efficiency – but enrollee risk differences do not. Therefore, there is no incentive to distort benefits to cream skim low-risk enrollees.

In reality, consumers vary not just in their medical risk but also in their *demand for the star hospital,* which is partly a function of preferences. Preferences for the star hospital influence both plan demand and costs, creating a positive correlation between ΔD_{ij} and ΔC_{ij} (conditional on risk) that leads to adverse selection. To formalize this, let $s_{i,h}(n_j)$ be a patient demand function that is i's probability of choosing h (under network n_j), which I assume for expositional simplicity is constant across diagnoses d. Suppose that the expected cost for treating diagnosis d is no longer a fixed ω_d but equals $\omega_d\tau_h$, where τ_h is a multiplier capturing the cost impact of provider h through both treatment intensity (quantity of care) and negotiated prices.[8] Under this richer "networks model," risk-adjusted costs equal:

$$\textbf{Networks model:} \quad C_{ij}^{RA}(n_j) = \underbrace{(R_i/\varphi_i)}_{\text{Unobserved risk}} \times \underbrace{\left[\sum_h \tau_h \cdot s_{i,h}(n_j)\right]}_{\text{Cost of chosen providers}\equiv\kappa_{ij}(n_j)} \tag{4}$$

where $\kappa_{ij}(n)$ is the (utilization-weighted) average cost of i's chosen providers, which also depends on the network. The equation shows the two dimensions of risk-adjusted cost heterogeneity: (1) unobserved medical risk, R_i/φ_i, and (2) the costliness of a patient's chosen providers, $\kappa_{ij}(n)$. The latter is likely to be particularly large for patients with high demand for the expensive star hospital.

Three reasons suggest that the cost heterogeneity in (4) is likely to create problems for standard

[8] My empirical work (Sections 5-6) weakens these assumptions, allowing patient choice probabilities to vary by diagnosis and for τ_h to vary by hospital-insurer pair. It also attempts to separate out τ_h into components occurring through treatment intensity vs. provider prices.

risk adjustment – even excellent risk adjustment that perfectly measures risk. First, heterogeneity due to use of expensive providers comes from varying patient demand (especially for the star hospital), which is partly a function of non-medical preferences. The variables entering risk adjustment typically do not include even observed determinants of preferences (e.g., distance), much less unobservable determinants; they are therefore unlikely to capture heterogeneity in $\kappa_{ij}(n_j)$.

Second, the same demand leading to high unobserved costs (via high $\kappa_{ij}(1)$) also affects plan choice. This creates a *direct* link between plan demand and unobserved costs, setting up the correlation between $C_{ij}^{RA}(1)$ and ΔD_{ij} that implies an adverse selection incentive in (2).[9]

Finally, and most fundamentally, costs due to varying provider choices are not exogenous (like medical risk) but *endogenous to the network*. Covering the star hospital affects costs by shifting provider choices, allowing patients to use the more expensive star hospital. Importantly, the incremental costs, ΔC_{ij}, are unlikely to be uniform across consumers. They will instead be higher for people with greater propensity to choose the star hospital. This is precisely the group likely to have high demand for a plan covering the star hospitals, setting up the conditions for "selection on moral hazard" (Einav et al., 2013). Selection on moral hazard poses a challenge for risk adjustment, since a single risk score φ_i (invariant to the network) cannot accurately capture independent variation in both $C_{ij}(0)$ and ΔC_{ij}. If regulators set risk scores based on $C_{ij}(0)$, they preserve selection on moral hazard. If they instead set risk scores based on $C_{ij}(1)$, they create an implicit subsidy for the broader network. While such a subsidy could be desirable, it is a policy with tradeoffs, not a mere technical fix. Fundamentally, selection on moral hazard complicates risk adjustment because it becomes tied up with fundamental economic cost-quality tradeoffs (Einav et al., 2016).

Adverse selection lowers the profitability of covering the star hospital. This may lead to standard implications: higher premiums (and thus lower enrollment) in star-hospital covering plans or in the extreme, full "unraveling" of star hospital coverage. In addition, adverse selection may influence hospital-insurer bargaining by *disciplining hospital market power*. Adverse selection effectively improves the insurer's bargaining threat point so may result in lower star hospital prices without exclusion (Ho and Lee, 2017). What occurs is an empirical question that will depend on the setting.

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.[10] Enrollees could choose among competing private plans in a

[9] The idea that demand for networks is driven by expected utilization is a core idea in the influential "option demand" model of Capps, Dranove and Satterthwaite (2003) – which is the foundation of most hospital-insurer bargaining models – but the implication for adverse selection has not been pointed out previously.

[10] 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.

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.[11] This structure lets me study plans that differ in network but are nearly identical on other dimensions.

Second, like the ACA, CommCare used sophisticated policies to address risk selection. In addition to benefit regulation and subsidies, it risk adjusted insurer payments.[12] Specifically, the exchange used demographics and past diagnoses to assign each enrollee i a "risk score" (φ_i) predicting their relative costliness. An insurer setting price P_j received revenue $\varphi_i P_j$ 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), CommCare's methods were state-of-the-art. The one notable limitation was the use of "prospective" risk scores based only on prior-year claims, whereas the ACA uses a "concurrent" risk score (the HHS-HCC method) based on current-year claims. While prospective risk adjustment limits incentive 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 prospective 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. *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 is the state's largest health system, giving it substantial market power. 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 known to drive high hospital prices: star status (Ho, 2009) and high market share (Cooper et al., 2019).

Table 1 shows these high Partners hospital prices in the CommCare data, drawing on estimates from the price model in Section 5.1. The table reports inpatient price estimates for the 10 highest-price hospitals in the data. Column (1) shows raw average payments per admission, and columns (2)-(4) report estimates of relative price and patient severity (vs. an average of 1.0 for each). The two star Partners hospitals (MGH and BWH) are the most expensive by a large margin, with relative prices of about 1.60, more than 20% above the next-highest hospital.[13] The star hospitals also treat some

[11] 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.

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

[13] A natural question is whether these high prices reflect high costs or markups. The answer appears to be both. Based on a state report of average cost per severity-adjusted admission (CHIA, CHIA), BWH 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 CommCare prices (since the casemix adjustment may differ), note that prices exceed the average by a larger percent (58-62%) than costs (30-50%), suggesting that markups 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	**Hospital Price Model** Relative Price Estimate	*Std. Err.*	Rel. Patient Severity
				(1)	(2)	(3)	(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. 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 inpatient price model (see Section 5.1), and column (4) reports average patient severity. Both prices and severities are relative measures, with a mean of 1.0 in the full data. (Price has mean 1.01 for the average hospital in this table because the sample is restricted to in-network admissions.)

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 the top-10 price list, and all six have above-average severity.

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 shifts in plan quality. 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, paying zero for all plans. Higher-income enrollees get the same plans but pay more on the margin for higher (pre-subsidy) price plans. This sets up a natural identification strategy for premium coefficients in my plan demand model. I discuss this strategy and the underlying premium variation further in Appendix B.

3.2 Administrative Data: Plan Enrollment and Insurer Claims

I use administrative data on enrollment and insurance claims for all CommCare plans and enrollees from fiscal 2007-2014.[14] For each (de-identified) enrollee, I observe demographics, plan enrollment history, and insurance claims. The claims include patient diagnoses, services provided, the provider identity, and actual amounts paid by the insurer. I use the raw data to construct the following three analysis datasets:

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 I observe networks. Constructing an admission-level dataset from the insurance claims – which often have multiple claims per admission – is an involved process; I discuss details in Appendix A.1. For each admission, I use the claims to observe the treating hospital, the principal diagnosis and diagnosis-related group (DRG), comorbidities, and total insurer payments (including both facility fees and physician professional payments). 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 zip code centroid to each hospital.[15] I use this dataset to estimate the hospital price and choice models.

Plan Choice and Cost Dataset The second dataset is for insurance plan choices and 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 enrollees can switch plans. These situations differ in their default outcomes: new and re-enrollees must actively choose a plan,[16] 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 data provider. Finally, I code the hospital or CHC (if any) at which the outpatient care was delivered

[14] 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.

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

[16] 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.

Figure 1: Changes for Network Health around 2012 Network Change

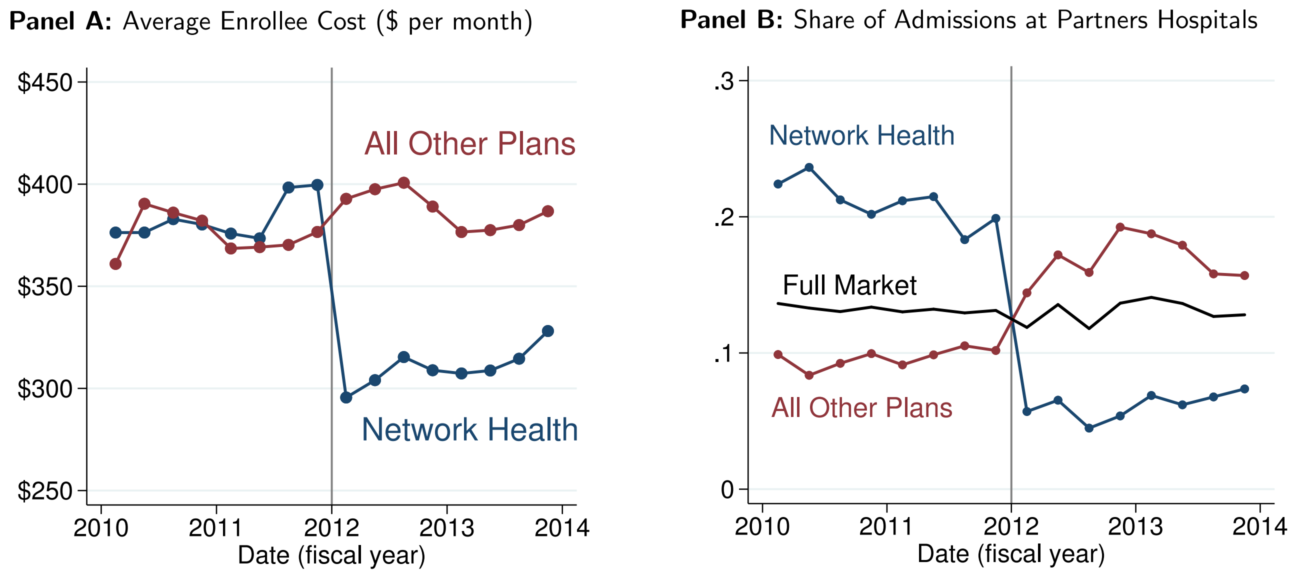

NOTE: The figures show average enrollee cost per month (left graph) and Partners hospital use shares (right graph) by enrollees in Network Health and all other CommCare plans. Each point is a quarterly average, and the vertical line marks the point where Network Health drops Partners from its network. 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.

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 A.1 reports summary statistics. The data include 624,443 unique enrollees making 1,684,203 plan choices and having 70,094 hospital admissions. The average age is 39.9, and 47% of enrollees are below-poverty so are fully subsidized. There is substantial flow 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 for plan demand estimation.

3.3 Star Hospital Coverage and 2012 Network Change

Plan hospital networks vary significantly, including in coverage of the star hospitals. Overall statistics on the size of hospital networks are reported in Appendix B.3. Here, I focus on the coverage of the star Partners hospitals. Up to 2011, three of the four Boston-area insurers covered the star Partners hospitals.[17] My empirical work exploits a major change in Partners coverage in fiscal 2012. In 2012,

[17] These three plans were 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

the exchange introduced new rules encouraging insurers to compete more aggressively on premiums.[18] In response, two plans (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 less prestigious hospitals.[19]

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 offsets 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. 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 changes in their networks. The two plans still covering Partners (CeltiCare and NHP) received over 90% of consumers who left Network Health in 2012, and their costs and Partners use rates rose sharply. 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.[20] 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 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 during fiscal 2013. NHP's status as the only plan to cover Partners has continued through at least 2019 in the state's post-ACA "ConnectorCare" program (the successor to CommCare).[21]

---

full Boston network.

[18]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.

[19]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.

[20]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."

[21]Moreover, NHP experienced significant financial challenges (e.g., losing \$100 million in 2014) and was forced to raise its prices substantially, leading its market share to fall into single digits by 2019. 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. 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.

4 Reduced Form Evidence of Adverse Selection

This section presents reduced form evidence of adverse selection on star hospital coverage consistent with the mechanisms in the theory in Section 2. To do so, I study the natural experiment created by Network Health dropping the star Partners hospitals in 2012, as just 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. Section 4.1 shows the main evidence from plan switching choices in 2012, and Section 4.2 examines the role of sickness and preferences in explaining switching choices.

4.1 Evidence from Plan Switching

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.[22] Changing plan choices come in two forms: (1) through plan switching by current enrollees and (2) through shifts in initial plan choices by new enrollees. My main analysis focuses on plan switching. This lets me study within-person demand changes and measure costs prior to the network change, to avoid conflating selection with causal effects of the network. The limitation is that plan switching is known to be affected by inertia. In robustness analyses, I examine new enrollee choices and find similar results (see Appendix C.1).

Figure 2 shows evidence of a large spike in switching out of Network Health in 2012, driven by consumers likely to have higher demand for Partners and other dropped hospitals. For the plan overall, the switching rate spikes to 11.3%, more than four times the 2.4% rate in 2010-11. 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. Switching rates spike to 22% for people within 5 miles of Partners, versus a steady 5% rate for those >25 miles away. Panel B shows that switching was even more concentrated among prior-year patients of the dropped hospitals (for outpatient care), a revealed preference indicator of demand. For prior-year Partners patients, the switching out rate spikes to 45% – a more than *twenty-fold* increase over the rate for the prior two years (2.1%). Switching also jumps to 24% for patients of other dropped hospitals (versus 1.7% in the prior two years). By contrast, switching for all other enrollees was much lower (3%) and essentially flat versus prior years.[23]

Figure 3 shows that 2012 switches were correlated with prior-year (2011) costs in a way consistent with adverse selection. Switchers out in 2012 represent a clear outlier in terms of high costs relative to other years when they have similar or lower costs than stayers. In 2012, switchers out have costs 108% higher than stayers ($675 vs. $324 per month). CommCare's risk adjustment narrows this

[22] The assumption throughout is that changing plan choices are caused by Network Health's narrower network and lower premium in 2012, and not other contemporaneous shocks. This assumption seems reasonable given the stability of other plans' networks at this time and given the pattern of switching I see in the data.

[23] 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.

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.

These patterns are consistent with the narrower network leading patients who value the excluded providers to switch plans to keep access to their preferred hospital or doctor.[24] Because these enrollees have high risk-adjusted costs (see Appendix Table A.4), this switching benefits Network Health via favorable selection. This story is quite intuitive. 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. High rates of plan switching occur despite the well-known fact of inertia in plan choice (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 may represent an important mechanism through which plan networks influence enrollee choices.

Robustness Checks In Appendix C.1, I implement three analyses to check the robustness of these adverse selection findings: (1) studying switching by zero-premium enrollees, for whom there is no concurrent change in Network Health's premium along with the narrower network; (2) examining new enrollee choices, which are not subject to inertia; and (3) showing similar evidence from CeltiCare's 2014 exclusion of Partners from its network. In all three cases, the findings discussed thus far appear robust. Switching patterns for zero-premium Network Health enrollees are similar to the main results. New enrollee demand for Network Health changes sharply at the start of 2012, with demand changes correlated with distance, prior Partners use, cost variables in ways similar to the main results. Finally, the evidence from CeltiCare's 2014 exclusion of Partners suggests similar response in plan switching, new enrollee choices, and adverse selection.

4.2 Role of Sickness and Preferences

The theory in Section 2 emphasizes two channels for costs and demand for the star hospitals: sickness and (non-medical) preferences. What role does each channel play in driving plan switching in 2012? To understand this, I run regressions to measure heterogeneity in the 2012 plan switching spike relative to prior years. Limiting the sample to current Network Health enrollees at the start of each year from 2009-12,[25] I estimate logit regressions of the form:

$$SwitchPlans_{i,t} = logit\left(\alpha + \beta \cdot X_{i,t} + \gamma \cdot X_{i,t} 1_{\{t=2012\}}\right) \tag{5}$$

where $X_{i,t}$ are enrollee characteristics (e.g., distance to Partners). In the regression, α and β capture plan switching patterns in 2009-11, and γ captures the *excess* switching in 2012.

[24]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).

[25]I exclude 2013-14 because the continued narrower network may affect switching patterns in those years. The results are qualitatively similar if those years are included, though excess switching rates are somewhat attenuated for the sickest enrollees, who continue switching out of Network Health at an elevated rate in 2013-14.

Figure 2: Spike in Switching Rates out of Network Health at 2012 Network Narrowing

Panel A: By Distance to Partners Hospital **Panel B:** By Prior-Year Patient Status

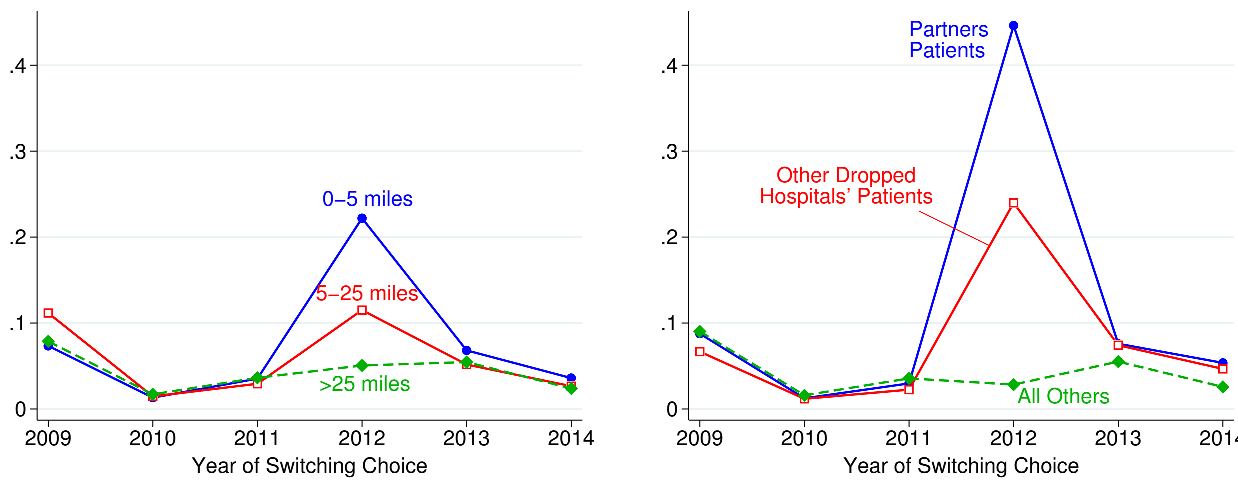

NOTE: These figures show switching out rates for Network Health enrollees at each year's open enrollment, separately by groups likely to correlate with demand for the providers dropped from network in 2012. Panel A shows that switching spikes for enrollees living closest to Partners hospitals. Panel B shows that switching spikes in 2012 for prior-year patients of Partners and other dropped hospitals (defined based on non-emergency room outpatient care).

Figure 4 plots estimates of excess switching odds ratios ($= \exp(\gamma)$) for $X_{i,t}$ variables capturing proxies for preferences (distance to the nearest Partners hospital) and sickness (risk score quantiles).[26] Each panel is a separate regression to ease interpretation.[27] The first panel shows that switching increases with proximity to Partners. People living >25 miles away switch at similar rates in 2012 as prior years (odds ratio = 1.11), but the switching odds spike rises to a 9.81-fold increase for people living within 2 miles. The second panel shows that switching also rises with sickness (captured by the HCC risk score), with an especially large increase for the sickest 5% of enrollees (odds ratio = 9.72). The third panel shows that a similar relationship holds for *unobserved* sickness, defined as the ratio of of the HCC risk score to the risk score used by CommCare.

The final panel shows that distance and sickness both matter conditional on each other, consistent with a model where both contribute to the utility function driving choice. Even among the healthiest 20% of enrollees, people living within 5 miles of Partners show a substantial switching spike in 2012 (odds ratio = 3.21). Likewise, conditional on distance, sicker enrollees are much more likely to switch plans. Even among people living >25 miles away, the sickest 5% of enrollees show a switching odds spike of 3.50. Consistent with the combined role of distance and sickness, the group by far most likely to switch plans are the sickest enrollees who also live nearby Partners (odds ratio = 25.34).

[26] For sickness, I use the HCC risk score for the prior year (e.g., 2011 value for 2012 switching choice), which avoids any reverse causality whereby switching could lead to a change in risk score. The HCC measure is a "concurrent" measure based on current-year claims (e.g., the 2011 score is based on 2011 claims), while CommCare's official risk score is based on prior-year claims.

[27] Results are similar with a multivariate regression; see Appendix Figure A.13. Appendix C.2 also shows that sickness and distance impact switching rates even conditional on prior-year patient status – i.e., estimating separate regressions for prior-year Partners patients and individuals who did not use a dropped hospital.

Figure 3: Adverse Selection Evidence: Average Cost of Switchers and Stayers ($ per month)

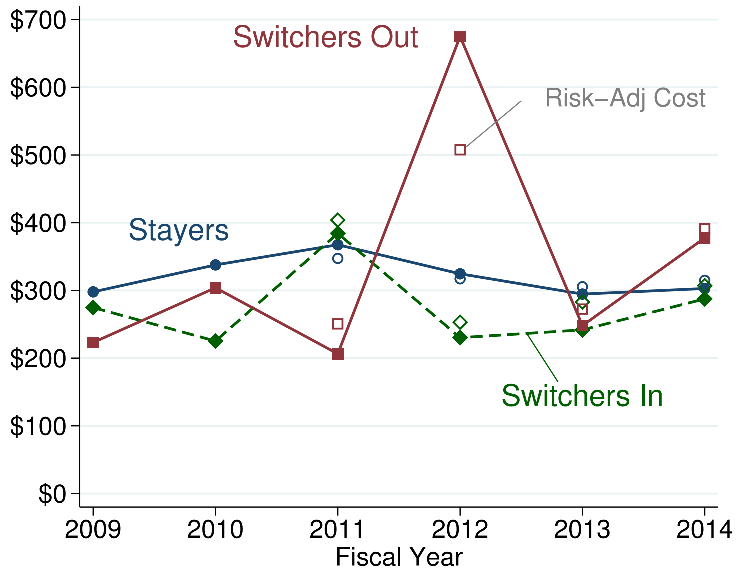

NOTE: The figure shows evidence that the 2012 plan switching spike shown in Figure 2 is consistent with adverse selection on star hospital coverage. The figure plots average prior-year costs of stayers, switchers out of, and switchers into Network Health by year. The connected series (with solid points) are raw average costs, and the open points are risk-adjusted costs using the exchange's method (which began in 2010, so is available for prior-year costs starting in 2011). The data show that 2012 is a clear outlier for selection patterns, with switchers out having much higher costs than stayers, and switchers in having lower costs.

These results suggest that *both* sickness and non-medical preferences drive demand for Partners coverage. A natural question, then, is how quantitatively important each factor is. To study this, I use a decomposition method of Shorrocks (2013) to quantify the role of sickness and preference covariates in explaining variation in two metrics of Partners demand: (1) switching plans in 2012, and (2) being a Partners patient in 2011. Appendix D.1 discusses the method details.[28]

Overall, the estimates suggests that (while both matter) distance – which is just one factor in preferences – is quantitatively more important than sickness. Even for the most detailed sickness specification, distance accounts for 56-69% of the explained variation in the demand metrics. "Observed risk" variables used in risk adjustment (age and the CommCare risk score) explain only 2-8% of variation, and a much richer set of "unobserved risk" measures derived from the claims explains another 28-35%. Moreover, there is substantial unexplained variation, which seems more likely to reflect unobserved preferences (which are hard to measure) than sickness (which is relatively well measured in the claims data). Consistent with a role for unobserved preferences, being a prior-year patient of Partners or another dropped hospital is by far the strongest predictor of switching plans, explaining more of the variation than all the sickness and distance measures combined. The large impact of having

[28]Briefly, the method calculates the contribution of each group of covariates to the (pseudo) R^2 of a logit regression of a Partners demand outcome on sickness and preference covariates. It accounts for complementarity among covariates by calculating the Shapley value – essentially averaging over the marginal contribution to R^2 for every possible covariate ordering. I include covariates for distance, "observed" sickness (age and CommCare risk score quantiles), and "unobserved" sickness (HCC risk score quantiles and other diagnosis and utilization variables). The specification includes up to 64 sickness variables with the goal of flexibly capturing risk beyond the measures used for risk adjustment.

an existing relationship with Partners raises the question of whether this affects demand because of persistent heterogeneity or state dependence (loyalty to one's current provider). Appendix D.2 discusses these two channels further and provides some evidence that both are involved.

Figure 4: Spike in Plan Switching in 2012: Role of Sickness and Preferences (distance)

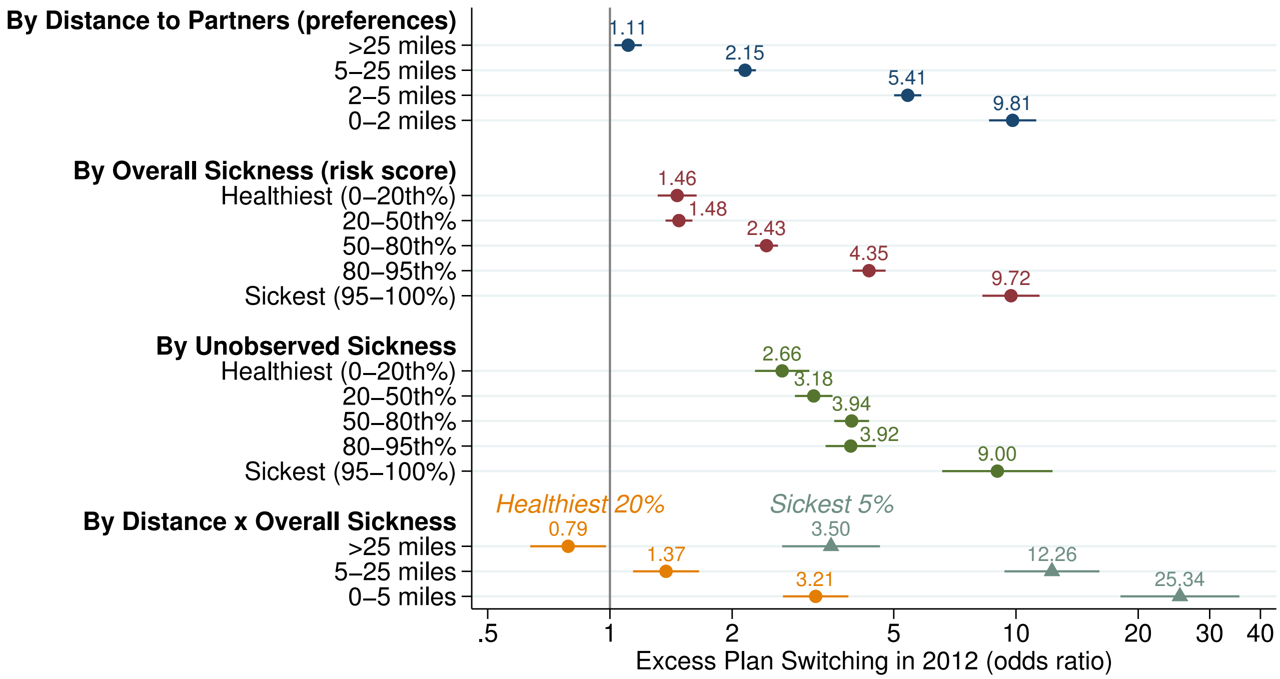

NOTE: The figure shows patterns of switching out of Network Health in 2012 for various sub-groups of enrollees, with odds ratios reported corresponding to estimates of $\exp(\gamma)$ from logit regression (5). The first panel shows patterns by enrollee distance to the nearest Partners hospital. The next panel shows patterns by overall sickness, defined as quantiles of the (prior-year) HCC risk score. The third panel shows patterns by unobserved sickness, defined as the ratio of of the HCC risk score (based on concurrently observed diagnoses) to CommCare's risk score (the retrospective measure used by the exchange). The final panel shows the interaction of distance and sickness, with estimates by distance for the healthiest 20% and sickest 5% of enrollees based on HCC risk score.

5 Understanding Costs Driving Adverse Selection

The evidence in Section 4 is consistent with a selection incentive to exclude the high-cost star hospital system. Doing so leads to reduced demand among consumers who value access to the star system, and who also have high risk-adjusted costs. This raises the question of *why* these consumers have high costs. What role is played by each of the two cost channels highlighted in the theory – greater medical risk and high costs due to use of expensive providers?

This section provides evidence on the role of these two cost channels. To do so, it uses two distinguishing features of the channels. First, risk should be reflected in higher *quantity* of care predictable by risk variables, while higher *prices* operate through the provider use channel. This motivates a cost decomposition into price vs. quantity in Sections 5.1-5.2. Second, the expensive provider use channel predicts *causal* cost reductions when the star hospitals are excluded, with larger reductions for groups more likely to use the star hospitals. Section 5.3 shows evidence of this prediction.

5.1 Decomposition of the Two Cost Dimensions

The theory in Section 2 shows that consumers may incur high costs through two dimensions: (1) medical risk (the standard channel) and (2) use of expensive providers (the non-standard channel). Equation (4) shows how costs can theoretically be separated into these two dimensions, as the product of medical risk (R_i) and the cost impact of chosen providers ($\kappa_{ij}(.)$). In equation (4), provider cost effects are given by a single factor τ_h, involving both prices and treatment intensities. In this section, I unpack the two, assuming that $\tau_h = \rho_{jh} \cdot \chi_h$, where ρ_{jh} is a negotiated price factor and χ_h is the hospital's treatment intensity (effect on quantity).

How can this decomposition be taken to the data? Start by noting that prices (ρ_{jh}) enter only through the provider choice/cost channel. Therefore, decomposing costs (C_{it}) into prices (P_{it}) vs. quantities (Q_{it}), which I discuss below, can begin to separate these channels, with price belonging to the second channel. Quantity variation, however, reflects a mix of both medical risk and provider treatment intensity.

Can these two be separated? To make progress, two observations are useful. First, medical risk reflects quantity that is *predictable based on patient risk factors* (e.g., age, diagnoses, risk scores), independently of the chosen provider. This motivates using a regression model to project quantity (Q_{it}) onto risk variables (Z_{it}) to capture "risk-predictable quantity," $\hat{Q}_{it}^{risk} = E(Q_{it}|Z_{it})$; see the method described below. The remaining "residual quantity," $\hat{Q}_{it}^{resid} \equiv Q_{it}/\hat{Q}_{it}^{risk}$, is ambiguous and may reflect either further unobserved medical risk or provider treatment intensity. One way to gain insight is to examine the relationship between residual quantity and the chosen provider, instrumenting for provider choice using distance to deal with sorting on unobserved risk. Second, note that the key distinction of the expensive provider use channel is that it is endogenous to the network (i.e., $\kappa_{ij}(n)$ varies with the network choice $n \in \{0, 1\}$), whereas quantity due to medical risk is fixed. This motivates examining (causal) changes in quantity and costs after the 2012 network change, which I do in Section 5.3 below.

To summarize, there are three ways of distinguishing the two cost dimensions:

1. Costs due to *high provider prices* reflect the expensive provider channel.

2. Quantity *predictable by patient medical variables* is medical risk. Residual quantity is ambiguous and may be a mixture of medical risk and provider treatment intensity, though we can gain insight by studying its relationship with the chosen provider.

3. *Causal changes in quantity and costs due to the network change* reflect the expensive provider use channel.

Sections 5.2 and 5.3 implement these three analyses. Before doing so, I provide an overview of the method for the price-quantity decomposition and estimating risk-predictable quantity.

Cost Decomposition Method I start by decomposing costs 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.[29] I define quantity as "price-standardized" utilization, or spending calculated at identical service-specific prices across providers. For each medical service $s \in \{1, ..., S\}$ (see definition below), define Q_s as the mean payment for s across all insurers and years. Price is defined as the (multiplicative) residual explaining observed insurer payments ($Paid_{a_{it},s}$) for each service instance (a_{it}) in the claims: $Paid_{a_{it},s} = Q_s \cdot P_{a_{it},s}$. This definition ensures that price is a relative measure centered around 1.0 for each service. Total quantity of care used by person i in year t equals: $Q_{i,t} = \sum_{a_{it} \in A_{it}} Q_{s(a_{it})}$, where A_{it} indexes the services used by the individual. The individual's average price of care (if $Q_{i,t} > 0$) is:

$$P_{i,t} \equiv \frac{C_{i,t}}{Q_{i,t}} = \sum_{a_{it} \in A_{it}} \left[\frac{Q_{s(a_{it})}}{Q_{it}} \right] \cdot P_{a_{it},s} \tag{6}$$

which is a quantity-weighted average price across all services an individual uses.[30] Applying this decomposition to the claims lets me calculate average quantity and price of care for each individual and for groups of enrollees, such as switchers vs. stayers in Network Health in 2012.

A key step in this method is defining the unit of medical services, s. I do so slightly differently for outpatient and inpatient care. For outpatient care, I use procedure codes (HCPCS codes), 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 (e.g., Cooper et al., 2019). Appendix E discusses details and shows descriptive statistics for the estimates.

After pulling out quantity, I project it onto medical risk observables (Z_{it}) to estimate "risk-predictable quantity." I do so using a two-part model, with a logit for the probability of positive quantity and log-linear regression for quantity conditional on positive. I output risk-predictable quantity as $\hat{Q}^{risk}_{it} = E\left[Q_{it}|Z_{it}\right] = f\left(Z_{it}; \hat{\theta}\right)$, where $f\left(.; \theta\right)$ is the two-part model's prediction function (see Appendix E.1). I implement this using two sets of Z_{it} variables: (1) only "observed risk" variables included in risk adjustment (age and CommCare's risk score), and (2) a broader set of variables from the claims (including diagnoses and the concurrent HCC risk score). After estimating $\hat{Q}^{risk}_{it}$, "residual

[29] The main excluded cost is prescription drugs. I exclude these 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. See Appendix E for further details.

[30] This price measure is 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 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.

Table 2: Decomposition of Switchers' High Costs: Price vs. Quantity

	Spending ($/month)	Quantity of Care				Provider Price Factor
		Overall Quantity	Predicted by Risk Vars.		Residual Factor	
			Used for Risk Adj.	All Risk Variables		
	(1)	(2)	(3)	(4)	(5)	(6)
A. Inpatient Care						
Stayers Mean	$47.7	$46.9	$58.0	$57.8	0.81	1.02
Switchers Mean	$152.8	$116.9	$71.3	$109.1	1.07	1.31
Ratio: Switchers / Stayers	*3.20*	*2.49*	*1.23*	*1.89*	*1.32*	*1.29*
Difference in Logs	1.16	0.91	0.21	0.64	0.28	0.25
	(0.11)	(0.10)	(0.02)	(0.05)	(0.08)	(0.02)
% of Log Diff. Explained	***100%***	***78%***	***18%***	***55%***	***24%***	***22%***
B. Outpatient Care						
Stayers Mean	$153.3	$161.7	$182.4	$197.2	0.82	0.95
Switchers Mean	$301.1	$309.8	$215.1	$282.4	1.10	0.97
Ratio: Switchers / Stayers	*1.96*	*1.92*	*1.18*	*1.43*	*1.34*	*1.03*
Difference in Logs	0.68	0.65	0.17	0.36	0.29	0.03
	(0.03)	(0.03)	(0.01)	(0.02)	(0.03)	(0.01)
% of Log Diff. Explained	***100%***	***96%***	***24%***	***53%***	***43%***	***4%***
C. Combined (IP + OP Care)						
Ratio: Switchers / Stayers	*2.26*	*2.05*	*1.19*	*1.54*	*1.33*	*1.10*
Difference in Logs	0.81	0.72	0.18	0.43	0.29	0.10
	(0.04)	(0.04)	(0.01)	(0.02)	(0.03)	(0.01)
% of Log Diff. Explained	***100%***	***88%***	***21%***	***53%***	***35%***	***12%***

Note: The table provides evidence on the source of costs driving adverse selection by decomposing cost differences between stayers and switchers out of the plan that narrows its network in 2012. All variables are for 2011 when both groups were in the same plan that covered the star hospitals. For the decomposition method, see Section 5.1. For inpatient costs (panel A), outpatient costs (panel B), and the sum of the two (panel C), the columns decompose switcher-stayer differences into components: overall quantity (col. 2), risk-predictable quantity (cols. 3-4); residual quantity (col. 5); and provider prices (col. 6). Columns 3 and 4 differ in the risk covariates used. Col. 3 includes only "observed risk" variables used in CommCare's risk adjustment: age groups and CommCare's risk score (entering with a flexible 11-part spline). Col. 4 adds concurrent variables observed in 2011: diagnoses and a spline in the HCC risk score.

quantity" is defined as the remaining factor explaining quantity: $\hat{Q}^{resid}_{it} \equiv Q_{it}/\hat{Q}^{risk}_{it}$.

Putting everything together, individual-level costs equal the product of three factors: $C_{it} = \hat{Q}^{risk}_{it} \cdot \hat{Q}^{resid}_{it} \cdot P_{it}$. This relationship also holds at a group level for (appropriately weighted) averages:[31]

$$\overline{C}_{g,t} = \overline{Q}^{risk}_{g,t} \times \overline{Q}^{resid}_{g,t} \times \overline{P}_{g,t} \tag{7}$$

This equation lets me decompose the share of group cost differences (e.g., stayers vs. switchers in 2012) that are driven by (1) risk-predictable quantity, (2) residual quantity, and (3) provider prices. Its multiplicative form suggests decomposing log differences for each factor, which are additive.

[31] $\overline{P}_{g,t}$ is average prices weighted by enrollee quantity (Q_{it}), and $\overline{Q}^{resid}_{g,t}$ is the average residual weighted by risk-predicted quantities ($\hat{Q}^{risk}_{it}$).

5.2 Cost Decomposition Results

I now apply the method just outlined to decompose cost differences between switchers out vs. stayers in Network Health in 2012 after it narrows its network. This sheds light on the source of cost differences correlated with demand for the excluded hospitals – in other words, the cost differences driving adverse selection. As in previous analyses, all outcomes and covariates are for 2011 when both groups were in the same plan and had access to the star hospitals.

Appendix Figures A.16-A.17 show descriptive plots of components of the decomposition for switchers vs. stayers, both overall and conditional on enrollee risk score. The patterns suggest that switchers are high-cost on nearly all metrics. Switchers have: (1) higher risk-predictable quantity using either "observed risk" factors (used by CommCare) or all risk measures, (2) higher residual quantity conditional on risk, and (3) higher inpatient prices, associated with greater use of the high-price Partners hospitals. These differences hold across the risk score distribution, suggesting that they are true for both sick and healthy. The lone exception for which there is little difference between switchers and stayers is outpatient prices, which I discuss further below.

Table 2 quantifies the contributions of each factor to switcher-stayer cost differences. Results are shown separately for inpatient care (panel A) and outpatient care (panel B), with the panel C showing the sum of the two. Spending (column 1) is substantially higher for switchers than stayers, by a factor of 3.20 (or +220%) for inpatient and 1.96 (or +96%) for outpatient costs. For the two combined, stayers have 126% higher costs – similar to the 108% excess for total costs (see Figure 2).

The remaining columns decompose the higher spending into price and quantity, with the bottom row of each panel showing the share of log differences explained by each. Three findings stand out. First, quantity explains the majority of switcher-stayer cost differences, with most quantity differences linked to medical risk. Using only "observed risk" factors used in risk adjustment (column 3) explains 18-24% of cost differences, while adding a broader set of risk variables explains 53-55% of differences (column 4). This large share shows that medical risk is still the main driver of selection, even in this setting where provider costs/choices matter. Moreover, it shows the value of better risk measurement. Comparing columns 3 vs. 4 indicates that "unobserved risk" – not captured by CommCare's risk adjustment but predictable using concurrent risk measures – explains 29-37% of adverse selection.

The second key finding is that provider prices (column 6) explain a meaningful 22% share of inpatient cost differences, though only 4% for outpatient costs. This indicates that the provider choice/cost channel matters for adverse selection.[32] The higher inpatient prices are entirely accounted for by switchers' >4x higher propensity to choose Partners hospitals (69% share for switchers vs. 15% for stayers), whose inpatient prices are 45% above-average. For outpatient care, switchers are also much more likely to choose Partners (33% vs. 6% share), but interestingly Partners' outpatient prices are not high (they are within 3% of the statewide average).

The third finding in Table 2 is that residual quantity (column 5) explains a substantial share of cost differences: 24% for inpatient and 43% for outpatient care. As noted, this component is more

[32]Moreover, I find that very little of the switcher-stayer price differences can be explained in regressions that control for medical risk observables, with the switcher-stayer ratio decreasing only from 1.29 to 1.25 when I control for risk. This is consistent with the findings in Section 4.2 that observable sickness explains only a small share of the demand for the star hospitals.

challenging to interpret. It may reflect either further unobserved risk *or* provider impacts on treatment intensity. To gain additional insight, Appendix Table A.9 examines how this residual quantity relates to propensity to use Partners, both using the raw OLS relationship and using distance as an instrumental variable. Partners patients consistently have high residual quantity, with levels about 20% higher in the IV specification using distance. This evidence, therefore, is consistent with both unobserved risk and provider effects contributing to residual quantity. To understand this further, I next analyze how costs change when the network is narrowed.

5.3 Evidence from Causal Cost Changes due to Narrower Network

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 examine whether dropping Partners has a *causal effect* on enrollee-level medical spending and how the effect varies across enrollees. In addition to testing this idea, these estimates are used in the 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 "stayers" continuously enrolled in Network Health from 2011-2012, relative to a control group of stayers in other plans. Limiting the sample to stayers and the 2011-12 period, I run a Poisson regression with individual and time fixed effects. The specification is:

$$E\left(C_{i,j,t}\right) = \exp\left(\alpha_i + \beta_t\left(Z_i\right) + \gamma\left(Z_i\right) \cdot 1_{\{j=NH, t\geq 2012\}}\right) \tag{8}$$

where $C_{i,j,t}$ is insurer cost on individual i at time t, α_i is an enrollee fixed effect, $\beta_t(.)$ 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 (8) is estimated by maximum likelihood (using "xtpoisson, fe" in Stata), with standard errors clustered at the i level. The coefficients of interest are $\gamma(Z_i)$, which capture the differential cost change for Network Health stayers in 2012. Note that (8) is analogous to standard difference-in-differences but in a non-linear model.[33] The implied (multiplicative) effect on costs equals $\exp(\hat{\gamma}(Z_i))$, and the percent change is $\exp(\hat{\gamma}(Z_i)) - 1$. I also estimate event study versions of (8) that allow $\gamma(.)$ to vary with time.

Figure 5 plots results from the event study version of (8), 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

[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.

Figure 5: Cost Reductions for Stayers after 2012 Network Change

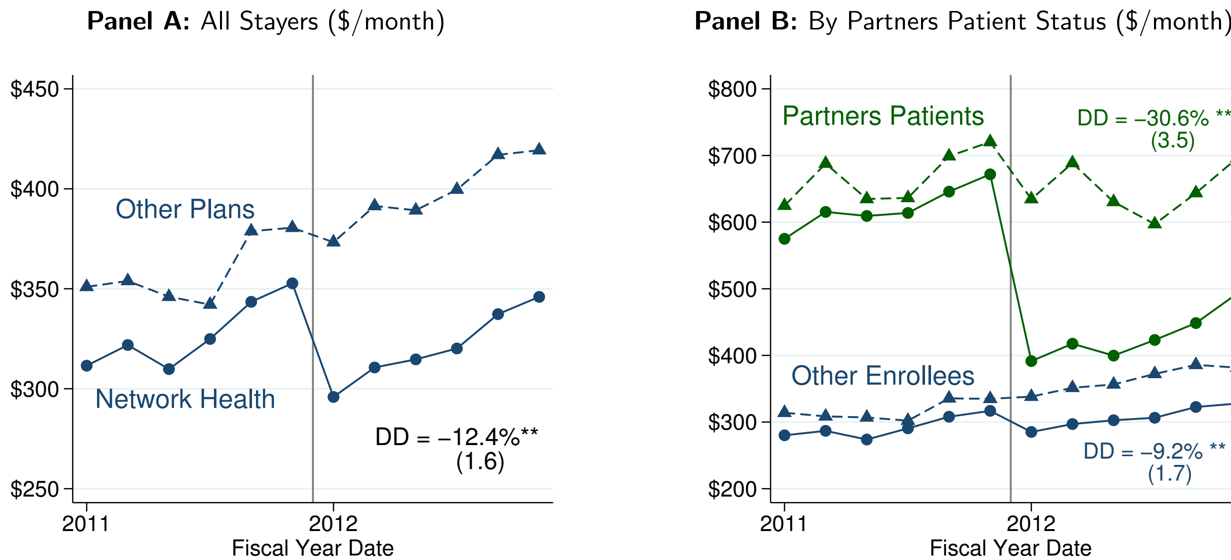

NOTE: These graphs show estimates from cost regressions with individual fixed effects corresponding to the event study version of equation (8). The sample is "stayers" continuously enrolled in Network Health or other plans between 2011 and 2012, when Network Health narrows its network. The outcome variable 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.

Health's costs aside from the one-time fall at the start of 2012.[34]

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 5 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 larger cost reductions at the start of 2012, both in levels (-$175 vs. -$30 per month) and in percentage terms (-30.6% versus -9.2%). Appendix Figure A.19 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).[35]

[34] Appendix Figure A.18 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.

[35] 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 5B). 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

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 A.13 (column (3)) confirms this finding in a richer specification of (8) that allow for richer Z_i heterogeneity on prior use, distance, medical risk factors, and demographics.

Appendix F.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 – comes 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 in the decomposition), where Partners prices are not high but they may deliver more intensive services.[36] These estimates provide further evidence that cost effect of using expensive providers involves both higher prices and treatment intensity, a finding also consistent with the evidence in Gruber and McKnight (2016).

6 Policy Analysis and Welfare Tradeoffs

6.1 Insurance Demand and WTP for Star Hospitals

To estimate consumers' valuations for star hospital coverage, I use the enrollment dataset to estimate a multinomial logit plan choice model. I treat individuals' timing of participation in the market as exogenous and model just their choices among plans.[37] Plan choices are made at two times: (1) new enrollments in the exchange (including re-enrollments after a break) and (2) plan switching decisions at annual open enrollment. For consumer i choosing at time t, the utility for plan j equals:

$$U_{i,j,t}^{Plan} = \underbrace{\alpha\left(Z_{it}\right) \cdot Prem_{i,j,t}}_{\text{Subsidized Premium}} + \underbrace{V\left(N_{j,t}; Z_{it}, \beta\right)}_{\text{Network Value}} + \underbrace{\delta\left(Z_{it}\right) \cdot 1\{CurrPlan_{i,j,t}\}}_{\text{Inertia (current enrollees)}} + \underbrace{\xi_{j,t}\left(Z_{it}\right)}_{\text{Plan dummies}} + \epsilon_{i,j,t}^{Plan} \qquad (9)$$

In addition to the "logit" error ($\epsilon_{i,j,t}^{Plan}$), plan utility depends on four plan characteristics: (1) subsidized premiums, (2) provider networks, (3) inertia for current enrollees in their current plan, and (4) unobserved quality, captured by a rich set of plan dummy variables. Because a key goal is to capture *heterogeneity* across consumers in price sensitivity and network valuation, I allow utility coefficients to vary with a rich set of consumer characteristics (Z_{it}), including income groups, age-sex groups, immigrant status, and deciles of the HCC risk score (plus an additional dummy for the top 5%). Appendix

further enrollees of 11.1% (s.e. = 1.8%); see Appendix Figure A.20 for event study estimates.

[36] 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 A.18 do not show much evidence that cost reductions diminish over time. But Figure A.19 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.

[37] The key assumption for my purposes is that plan network changes lead consumers to switch plans but do not affect exchange participation. This 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 premium of the cheapest plan after subsidies – the main variable likely to affect exchange participation – is set directly by the exchange's (price-linked) subsidy rules and does not change if insurers reduce premiums. To assess this assumption, Appendix Figure A.11 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.

F.3 lists the detailed interaction terms for each covariate and shows estimates. I now describe more detail about the four plan characteristics in the model:

- **Subsidized premiums** are observed and included directly. Premiums vary for two reasons: (1) because of insurer pricing, which occurs at the plan-year level, (in some years, separately across five regions) and (2) because of subsidies, which vary across five income groups. As discussed below, I setup the econometrics to identify premium coefficients only from variation due to subsidies by including plan dummies that soak up all variation due to insurer pricing.

- **Provider networks** are 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 past work (starting with Capps et al., 2003) by including a "network utility" measure derived from an estimated hospital choice model. Appendix F.1-F.2 present the model estimates and construction of network utility. Second, I include variables for whether a plan covers hospitals with which the consumer has past outpatient relationships (or the share covered if there are multiple). I 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.

- **Inertia (for current enrollees)** is well known to affect health insurance choices (e.g., Handel, 2013; Ericson, 2014).[38] To capture inertia in a simple way, I include a dummy for current enrollees' current plan, with coefficients $\delta(Z_{it})$ that vary with observables. This ensures that the model matches 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 with only new/re-enrollees for whom inertia is not relevant.

- **Plan dummy variables** are included both to capture unobserved plan quality (e.g., insurer reputation; see Starc, 2014) and to aid in identification of premium coefficients. I include separate plan dummies by region-income group ($\xi_{j,Reg,Inc}$) and region-year ($\xi_{j,Reg,Yr}$), as well as plan interactions with age-sex groups and risk score quantiles to allow variation with medical risk.

Identification of Premium Coefficients Properly identifying premium coefficients requires isolating variation orthogonal to unobserved plan quality/demand shocks. Rather than using instruments, I follow an alternate approach (see e.g., Nevo, 2000) of including detailed plan dummies to soak up all premium variation due to (likely strategic) insurer pricing so that remaining variation comes only from plausibly exogenous subsidies.

The logic works as follows. Insurers set (pre-subsidy) plan prices at either the region-year level (prior to 2011) or at the yearly level (2011+). Insurers by rule may not vary prices at a more detailed

[38] Inertia (or switching costs) may arise for a variety of underlying reasons. The most natural mechanism in the CommCare setting is an attention or hassle cost of switching plans, since current enrollees remain with their existing plan by default. Other possible reasons include an information/search cost of learning about other plans and real costs of switching plans (e.g., paperwork, or the costs of switching doctors). Because benefits in CommCare are standardized and I model provider networks directly, the latter explanation seems less likely to apply.

level than this. To avoid using this pricing for identification, utility includes plan-region-year dummies. This ensures identification comes only from premium variation across consumers *within a plan-region-year cell.*

This remaining variation comes only from subsidies. As Appendices B.1-B.2 detail, CommCare uses a complex subsidy schedule that creates variation by income (within a plan-region-year cell) in both premium levels and cross-plan differences. Notably, subsidies make *all plans free* for enrollees with incomes below poverty, while above-poverty enrollees pay higher premiums for higher-price plans. This structure makes demand patterns among below-poverty enrollees – who do not pay premiums – a natural "control group" for picking up shifts in unobserved plan quality.[39] To account for any persistent preference differences across income groups, utility includes plan-region-income group dummies. Therefore, premium coefficients are estimated from *differential premium changes* by income group for a given plan in a given region.[40]

This identification strategy is analogous to difference-in-differences (DD) in a non-linear model. As in standard DD, I include fixed effects to absorb all premium variation driven by endogenous factors, leaving only the exogenous (subsidy-driven) variation for identification. The assumption is that there are no further income group-specific demand trends/shocks (for a given plan in a given region) – i.e., no $\xi_{j,Reg,Inc,Yr}$ – that are correlated with premium changes. One simple test of this assumption is to examine whether demand trends are parallel between "treatment" (above-poverty) and "control" (below-poverty) groups around premium changes. Appendix Figure A.23 shows such a test, finding that monthly market shares are flat and parallel for treatment and control groups at all times except for the treatment group at the expected time (when premiums change at the start of each year).

Demand Estimates All variables entering the plan choice model are observed, so I estimate it by maximum likelihood. Appendix Tables A.11-A.12 show the estimates. Focusing on the main summary coefficients reported in Appendix Table A.11, column (2) reports the main specification including all enrollees.[41] 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 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.[42] There is substantial heterogeneity in price sensitivity, with less negative premium coefficients for higher-income, sicker, and older individuals.

[39] 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.

[40] In particular, a major source of identification is how market shares change for above-poverty enrollees when premiums increase/decrease, compared to changes in shares for the same plan among below-poverty enrollees. Appendix B.2 illustrates the logic by walking through an example, following the evolution in premiums for Network Health in a specific region (Boston) from 2010-2013.

[41] 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 distinguishing inertia vs. unobserved preferences. I therefore use column (2) for the remainder of the analysis.

[42] 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).

There is also substantial inertia in consumers' plan switching decisions, with the average coefficient of 4.413 (s.e. = 0.007).[43] 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 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 (Figure 2B), this coefficient is consistent with consumers placing a special value on star providers.

The estimates show substantial heterogeneity in network valuation via the interaction terms. Older, sicker, and higher-income enrollees have higher utility of networks covering their desired providers. In combination with these groups' smaller price coefficients, this implies higher willingness to pay for provider coverage. I analyze this heterogeneity and how it relates to costs 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 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.

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 $C_{ijt}^{obs}\left(1\right)$). Assume that realized costs equal expected costs ($C_{ijt}\left(1\right)$) plus an idiosyncratic shock: $C_{ijt}^{obs}\left(1\right) = C_{ijt}\left(1\right) + \epsilon_{ijt}$. If the variables defining group G are known at the time when expected costs are defined,[44] then $E\left(\epsilon_{ijt}|i \in G\right) = 0$ and expected costs for group G can

[43] 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).

[44] 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.

be estimated as the average of realized costs: $\overline{C}_{G,t}(1) \equiv \frac{1}{N_G}\sum_{i\in G} C^{obs}_{ijt}(1) \rightarrow E(C_{ijt}(1)|i \in G)$ 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 a consumer's *incremental cost* of the broader network, or $dC_i = C_{ijt}(1)/C_{ijt}(0)$. To do so, I draw on the causal estimates of Section 5.3, which are identified from 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 5 shows evidence of parallel pre-trends. I use the estimates of Poisson regression (8) 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 $d\hat{C}(Z_i) \equiv \exp(-\hat{\gamma}(Z_i))$, with the negative sign because $\gamma(.)$ comes from the reverse experiment of a narrower network. Appendix Table A.13 shows the results, with columns (3)-(6) reporting estimates for insurer cost, quantity, and prices. Given an estimate of either $\overline{C}_{G,t}(1)$ or $\overline{C}_{G,t}(0)$ from the data and $d\hat{C}(Z_i)$ from the regressions, I construct costs under the counterfactual network by multiplying/dividing each individual's observed costs by $d\hat{C}(Z_i)$ as appropriate.

A limitation of this method is that it infers incremental costs from stayers, who are a selected group. This raises two concerns. The first is whether the estimates of $d\hat{C}(Z_i)$ are *internally valid* estimates for stayers. I discuss and make the case for this in Section 5.3 above. The second is whether the estimates from stayers are *externally valid* when extrapolated to switchers with the same observables Z_i. This is more difficult to test, since I never observe switchers under the narrow network. The logic of selection on moral hazard suggests that $d\hat{C}(Z_i)$ might be unobservably larger for switchers, who are selected on high demand for the star hospitals. To the extent true, my estimates would be a conservative *under-estimate* of $\Delta Cost$, which would reinforce the finding that these are larger than consumer WTP for star hospital coverage.

6.3 Role of Selection Incentive 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 (2) in the theory. Table 3 shows the analysis. For simplicity I implement it on a balanced panel of enrollees in the CommCare market from the final quarter of 2011 to the first quarter of 2012.[45] Columns (1)-(2) show Network Health's premium, demand, costs, and profits for these two periods.[46] The next columns follow equation (2) in breaking down the profit change into selection incentives (columns 3-4)

[45] I do not include market exiters (leave during 2011) and new enrollees (join in 2012) because it requires more assumptions about their counterfactual plan choices under one of the networks. In practice for a range of assumptions, exiters and new enrollees appear to strengthen selection incentives and the profitability of the narrower network, suggesting that the results in Table 3 are conservative.

[46] 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 A.1. 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).

and price/cost changes with fixed enrollment (columns 5-6).

Table 3: Analysis of Incentives for 2012 Network Narrowing

	Observed Values		**Breakdown of Change in Profit**			
			Selection Incentive		**Fixed Enrollment**	
	2011 Broad Network	2012 Narrow Network	Switch Out	Switch In	ΔPrice	ΔCost
	(1)	(2)	(3)	(4)	(5)	(6)
Premium and Medical Costs ($/month)						
Plan Premium *(P)*	$423	$360	$423	$423	*-$63*	---
Raw Costs	$391	$289	$772	$223	---	*-$45*
Risk-Adjusted (AC^{RA})	$369	$285	$573	$260	---	*-$44*
Breakdown: $C^{RA}(0)$	$315	$285	$449	$228		
ΔC^{RA}	$54	$49	$124	$32		
Demand & Profits						
Demand *(risk-scaled)*	44,444	40,843	*-6,351*	*2,746*	40,843	40,843
Margin $(= P - AC^{RA})$	$54	$75	-$151	$163	*-$63*	*$44*
Total Profit ($million)	**$2.40**	**$3.05**				
ΔProfit 2011-12		***+$0.65***	***+$0.96***	***+$0.45***	***-$2.56***	***+$1.80***

Note: The table breaks down the profitability of Network Health's network narrowing in 2012. It implements equation (2) from the theory to decompose the change in profits (columns 1-2) into selection incentives (columns 3-4) vs. fixed enrollment price/cost changes (col. 5-6). Outcomes are measured in the final quarter of 2011 and first quarter of 2012, and the sample is restricted to a balanced panel of continuing enrollees in the market in both periods. Columns 1-2 show premiums, costs, demand, and profits (before admin costs) directly from the data. Columns 3-4 show the selection incentive, equal to the profitability of switchers $(P_j(1) - C_{ij}^{RA}(1))$ times their change in demand (ΔD_{ij}) from 2011-12. Columns 5-6 show the fixed enrollment price/cost changes $((\Delta P_j - \Delta C_{ij})\, D_j(0))$. Columns 3-6 use observed values where available or predictions from the cost model when not (e.g., costs of switcher in under the broad network). The rows in gray break down risk-adjusted costs into "baseline" costs under the narrow network $(C^{RA}(0))$ and incremental costs of the broader network (ΔC^{RA}).

The results illustrate both the strong overall incentive for a narrower network and the role of selection. Even though the plan cuts its premium substantially (by $63 per month), average costs fall by an even larger $102 (or 26%) in raw terms, and by $83 after risk adjustment (21%). Therefore, its profit margin increases by $21 per month (38%), outweighing a modest decline in demand and leading to $0.65 million higher profits.

Columns 3-4 show the large role of adverse selection in these changes, corresponding to the profitability of switchers in/out of the plan. The very high risk-adjusted costs of switchers out implies that the plan lost money on these enrollees (a margin of -$151 per month); their leaving the plan implied almost $1 million higher profits. Similarly, the low costs of switchers in implies high profitability (margin of +$163); their joining the plan increases profits by $0.45 million. Together, the selection incentive equals $1.41 million. This is about 60% of baseline 2011 profits, and 78% of the $1.8 million causal cost savings with fixed enrollment (column 6). Had there not been adverse selection, the plan would have lost money on this fixed set of enrollees, since the revenue losses from lower prices (col. 5) exceed the cost savings (col. 6).

The table also illustrates the interaction of selection and moral hazard, as suggested by Section 5.3. The high risk-adjusted costs of switchers out ($573 per month) reflects both high "baseline" cost under the narrow network ($C^{RA}(0) = \$449$) and a large incremental cost ($\Delta C = \$124$, or 28%). By contrast switchers in have both lower baseline (\$228) and incremental costs (\$32, or 14%). This pattern of selection on moral hazard contributes to the challenges of risk adjustment and the difficult welfare/policy tradeoffs involved, which I discuss next.

6.4 Analysis: WTP and Cost Curves for 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 (2010b, "EFC").[47] This approach provides a useful way of summarizing demand/cost primitives to understand the forces driving adverse selection and welfare. It works by ranking consumers in terms of decreasing WTP types for the broader network (call this ranking $s \in [0, 1]$) 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 (9):

$$\Delta WTP_i \equiv \frac{1}{-\alpha(Z_i)} \cdot [V(N_{NH,2011,}; Z_i, \beta) - V(N_{NH,2012}; Z_i, \beta)] \quad (10)$$

where $V(.)$ is the consumer's network valuation for the 2011 and 2012 network, converted into money terms by dividing by -1 times the premium coefficient.[48] The other key variables are costs, which are estimated using the cost model from Section 6.2. I plot cost variables 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. For simplicity, I focus on Network Health enrollees in 2011; results are similar if I examine other groups such as enrollees in 2012.

Figure 6 shows results. Panel A shows cost curves. Two results stand out. First, cost curves under the broad network (both raw and risk-adjusted) slope steeply downward with WTP, indicating strong adverse selection. Risk adjustment (light-blue dashed curve) makes a 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.

Figure 6B shows this result directly and plots the key curves for a standard welfare analysis: $\Delta WTP(s)$ and incremental costs, $\Delta Cost(s) \equiv \overline{C}(1; s) - \overline{C}(0; s)$. 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, the broader network with the star hospitals is

[47] 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).

[48] 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.

Figure 6: Cost and Willingness to Pay Curves for Broader Network

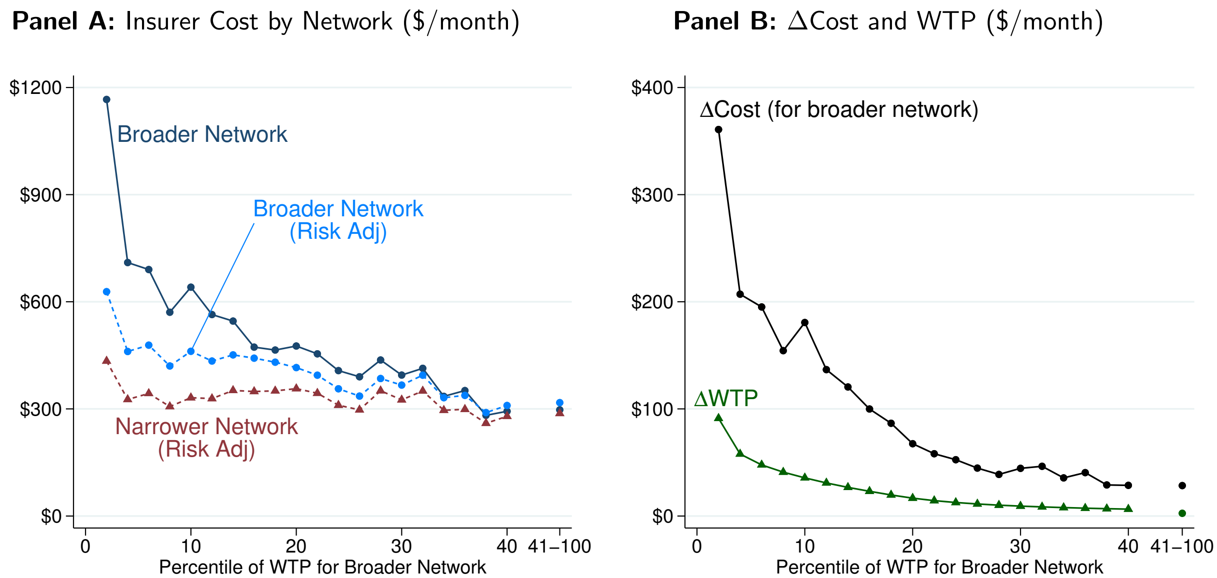

NOTE: These graphs show cost and willingness to pay (WTP) curves derived from the structural model estimates. The x-axis for both panels is the WTP type (s), the percentile ranking of WTP for Network Health's broader 2011 network that includes the star Partners hospitals, relative to the narrower 2012 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).

inefficient, as insurer costs exceed consumer value. 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[49] – their incremental costs are even larger (\$361 per month). Indeed, because $\Delta Cost$ is *steeper* than ΔWTP, social surplus ($= \Delta WTP - \Delta Cost$) is actually most negative for the highest-WTP types, consistent with the "backward sorting" pattern found by Marone and Sabety (2021). 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.[50] Nonetheless, the basic finding that costs of star hospital coverage are larger than consumers'

[49]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.

[50]A related issue is that WTP may diverge from consumers' true long-run value of star hospital coverage due to either behavioral biases or state dependent preferences. If some consumers are inattentive to networks when choosing plans,

WTP appears fairly robust. Appendix F.5 presents robustness checks on ΔWTP and $\Delta Cost$ estimates, including: (1) counting only quantity reductions in $\Delta Cost$, (2) recalculating $\Delta Cost$ using 10-50% lower Partners prices (reflecting possible markups above true marginal costs), (3) redefining ΔWTP based on a lower social marginal utility of money, and (4) counting in $\Delta Cost$ only savings from shifting to lower-price hospitals for inpatient care, as predicted by the hospital choice model.

For the first three analyses, the main result of $\Delta Cost > \Delta WTP$ continues to hold across the entire distribution. However, if cost changes occur only via inpatient care (analysis #4), $\Delta Cost$ is much smaller and now falls below ΔWTP. This suggests that consumers would be willing to pay for star hospital coverage if the only source of higher costs were shifting inpatient care toward higher-price hospitals. Consumers, however, are not willing to pay the much larger incremental costs that occur through higher quantity, especially for outpatient care. An important issue for future research is to better understand these quantity changes and whether they are clinically appropriate (but undervalued by consumers) or whether they reflect wasteful over-use.

7 Conclusion

As the use of market-based health insurance rises, an important question is how well 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 not only because they are sicker (the standard channel) but also 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 accepting lower negotiated 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 op-

the ΔWTP curve would be understated. State dependence (e.g., due to a cost of switching doctors), which I analyze further in Appendix D.2, has a more complicated impact. Normally, switching costs imply that short-run utility losses are *larger than* long-run losses, which reinforces the finding that WTP falls short of costs. But in this setting, $\Delta Cost$ is *also* driven by preferences for using star hospitals, and in the long-run a patient who loses Partners access and switches doctors will also have lower $\Delta Cost$ from regaining access to Partners. Thus, the long-run impact of state dependence on $\Delta WTP - \Delta Cost$ is ambiguous.

tions – 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. This poses a challenge for standard risk adjustment (Einav et al., 2016) and can make consumer sorting inefficient with any single pooled premium (Bundorf et al., 2012; Marone and Sabety, 2021). As a result, subsidies for generous coverage may not improve welfare.

The results suggest the importance of studying alternate policies to address these inefficiencies. Fundamentally, these problems go back to a basic sorting challenge: which patients should get access to expensive star hospitals? In the current system, consumers get access to star hospitals via their plan choice, after which the extra cost of these providers is largely covered 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 (Prager, 2020) or incentives for doctors to refer patients more efficiently (Ho and Pakes, 2014) – 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

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

Breyer, Friedrich, M Kate Bundorf, and Mark V Pauly, "Health care spending risk, health insurance, and payment to health plans," in "Handbook of health economics," Vol. 2, Elsevier, 2011, pp. 691–762.

Brot-Goldberg, Zarek C, 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*, 2017, *132* (3), 1261–1318.

Brown, Jason, 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*, 2014, *104* (10), 3335–64.

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

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

Carey, Colleen, "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D," *American Economic Journal: Economic Policy*, 2017, *9* (1), 38–73.

(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/ 2014.

_ , "Massachusetts Hospital Profiles: Data through Fiscal Year 2012," Technical Report, https://www.chiamass.gov/hospital-profiles/ 2014.

Chiappori, Pierre-André and Bernard Salanie, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 2000, *108* (1), 56–78.

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

Coakley, Martha, "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.

Cooper, Zack, 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*, 2019, *134* (1), 51–107.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya, "Can health insurance competition work? evidence from medicare advantage," *Journal of Political Economy*, 2021, *129* (2), 570–606.

Decarolis, Francesco and Andrea Guglielmo, "Insurers response to selection risk: Evidence from Medicare enrollment reforms," *Journal of Health Economics*, 2017, *56*, 383–396.

Einav, Liran, Amy Finkelstein, and Jonathan Levin, "Beyond Testing: Empirical Models of Insurance Markets," *Annual Review of Economics*, 2010, *2(1)*, 311–336.

_ , _ , and Mark R Cullen, "Estimating welfare in insurance markets using variation in prices," *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.

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

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

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

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

— **and Amanda Starc**, "Pricing regulation and imperfect competition on the massachusetts health insurance exchange," *Review of Economics and Statistics*, 2015, *97* (3), 667–682.

— **and** — , "How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics*, 2016, *50*, 71–85.

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

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

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

Gaynor, M and WB Vogt, "Competition among Hospitals.," *The Rand Journal of Economics*, 2003, *34* (4), 764–785.

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

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

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

Ghili, Soheil, "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance," *Marketing Science*, Forthcoming.

Glied, Sherry, "Managed care," in "Handbook of health economics," Vol. 1, Elsevier, 2000, pp. 707–753.

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

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

Health Care Cost Institute, "2014 Health Care Cost and Utilization Report: Analytic Methodology," 2015.

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

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

— **and** — , "Equilibrium provider networks: Bargaining and exclusion in health care markets," *American Economic Review*, 2019, *109* (2), 473–522.

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

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

Lavetti, Kurt and Kosali Simon, "Strategic formulary design in Medicare Part D plans," *American Economic Journal: Economic Policy*, 2018, *10* (3), 154–92.

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

Marone, Victoria R and Adrienne Sabety, "When Should There be Vertical Choice in Health Insurance Markets?," Technical Report, Working Paper 2021.

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

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

— **, 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*, 2015, *1* (1), 1–26.

Prager, Elena, "Health Care Demand Under Simple Prices: Evidence From Tiered Hospital Networks," *American Economic Journal: Applied Economics*, 2020, *12* (4), 196–223.

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

Sabety, Adrienne, "The Value of Relationships in Health Care," Technical Report, Working Paper 2020.

Shepard, Mark and Myles Wagner, "Reducing Ordeals through Automatic Enrollment: Evidence from a Health Insurance Exchange," Technical Report, Working paper 2021.

Shorrocks, Anthony F, "Inequality decomposition by factor components," *Econometrica*, 1982, pp. 193–211.

— , "Decomposition procedures for distributional analysis: a unified framework based on the Shapley value," *Journal of Economic Inequality*, 2013, *11* (1), 99.

Song, Zirui, 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*, 2012, *31* (8), 1885–1894.

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

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

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

Online Appendix:

Hospital Network Competition and Adverse Selection

Mark Shepard

A Appendix: 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[51] plus either a valid revenue code for "room and board" services[52] 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).[53] 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.[54] Using this episode sample, I then add on *all* claims (including professional and ancillary services) that occurred on a day the patient was admitted.[55] I also include emergency department

[51] 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.

[52] 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.

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

[54] 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.

[55] 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.[56]

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.[57] 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).[58] 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).[59] 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

[56]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.

[57]The vast majority (about 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.

[58]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.

[59]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 (about 1%), APR-DRG (about 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.[60] 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.[61]

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 C.1). However, I do not use it for the plan choice model or cost model estimation because I

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

[61]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 (2021) 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 A.1B.

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 A.11) differ slightly from those reported in Table A.1B because the plan choice model drops 3.5% of instances where individuals have only a single plan available.

Table A.1: Summary Statistics

Panel A: Hospitalization Dataset

Patient Characteristics

Variable		Mean
No. of Hospitalizations		70,094
Age		44.7
Male		48%
Emergency Department		65%
Principal Diagnosis	Mental Illness	14.9%
	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 A.11 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.

B Appendix: 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 A.1A 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.[62] Average enrollee premiums are shown in Panel B of Figure A.1, 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 A.1A 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

[62] 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 A.2 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.[63]

The plot shows how changes in Network Health's (and its competitors') pre-subsidy prices (Figure A.1A) 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 A.2 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 A.2. First, the specification for plan utility (equation (9) in the text) includes plan-region-year dummies (ξ_{j,Reg_i,Yr_t}) that absorb variation due to insurer pricing (which occurs at the plan-region-year or plan-year level) and in particular, 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 set at the start of every fiscal year

[63] 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 A.1: CommCare Plan Prices and Enrollee Premiums

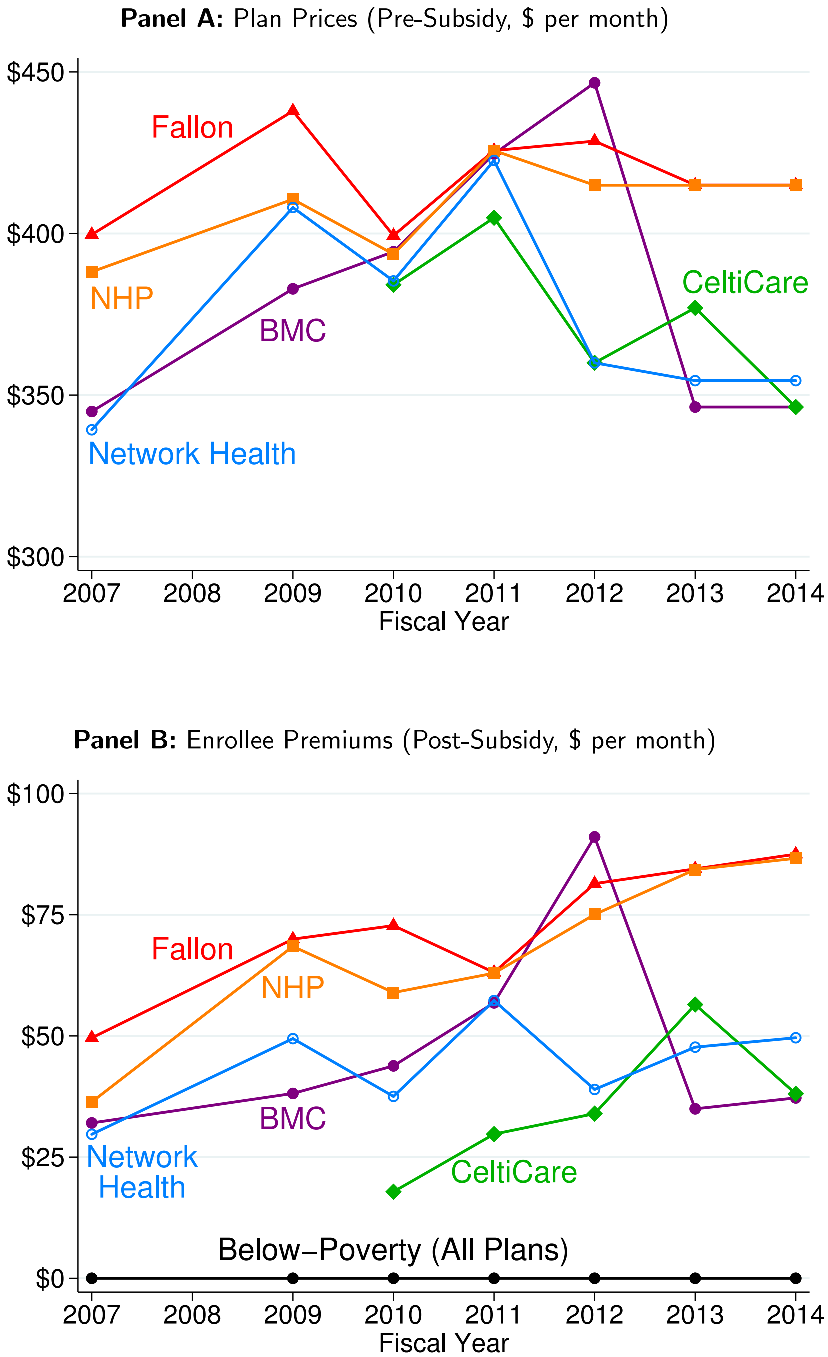

NOTE: The graphs show average pre-subsidy insurer prices (Panel A) and post-subsidy enrollee premiums (Panel B) for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying 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 A.2: 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 A.2 shows the distribution of relative premium changes for a plan between adjacent years, separately for premium decreases and increases (following the presentation in Figure A.23). 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 A.3 shows information on plans' share of hospitals covered (weighted by hospital beds), and Table A.3 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 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

Figure A.2: Identifying Premium Variation Example: Network Health (Boston region), 2010-13

Panel A: Enrollee Premium Levels by Income ($/month)

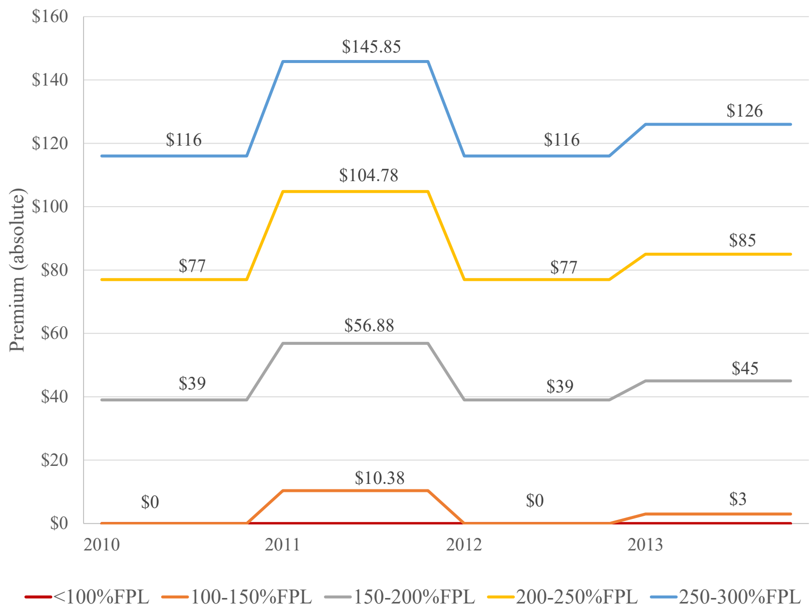

Panel B: Enrollee Premiums Relative to Cheapest Plan ($/month)

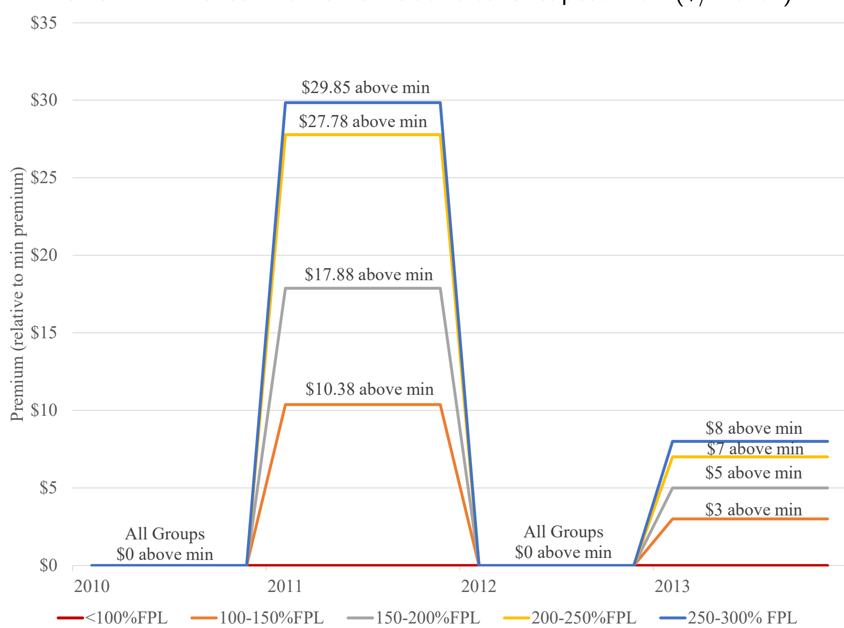

NOTE: The graphs shows the example of Network Health's (post-subsidy) enrollee premiums by income group over the 2010-2013 CommCare years. "FPL" refers to the federal poverty level. Pre-subsidy prices (and enrollee premiums) vary at the regional level in 2010, and the graph shows premiums specifically 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.

Figure A.3: Hospital Coverage in Massachusetts Exchange Plans

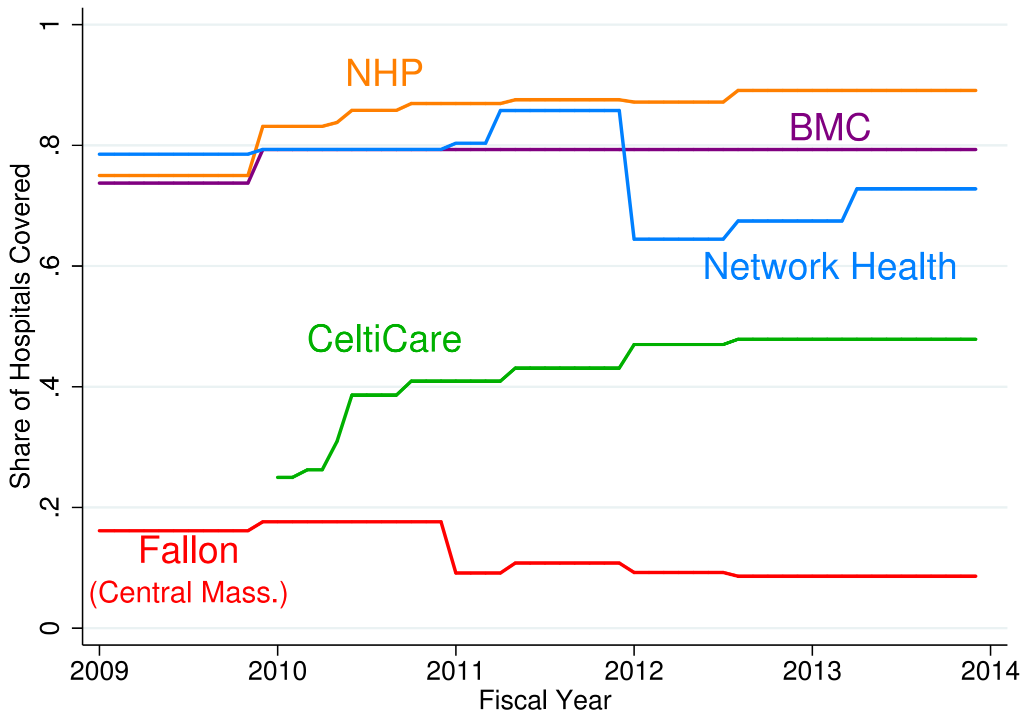

NOTE: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a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

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 A.3 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.

Table A.3: 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.

C Appendix: Robustness and Additional Analyses

C.1 Robustness Analyses on Adverse Selection Findings

The evidence in the body text (Section 4) focuses 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: (1) studying switching by zero-premium enrollees, for whom there is no concurrent change in Network Health's premium that could affect results; (2) examining new enrollee choices, which are not subject to inertia; and (3) showing similar evidence from CeltiCare's 2014 exclusion of Partners from its network.

(1) 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 A.4-A.5 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.

Figure A.4: Plan Switching and Selection for Network Health: Zero-Premium Enrollees

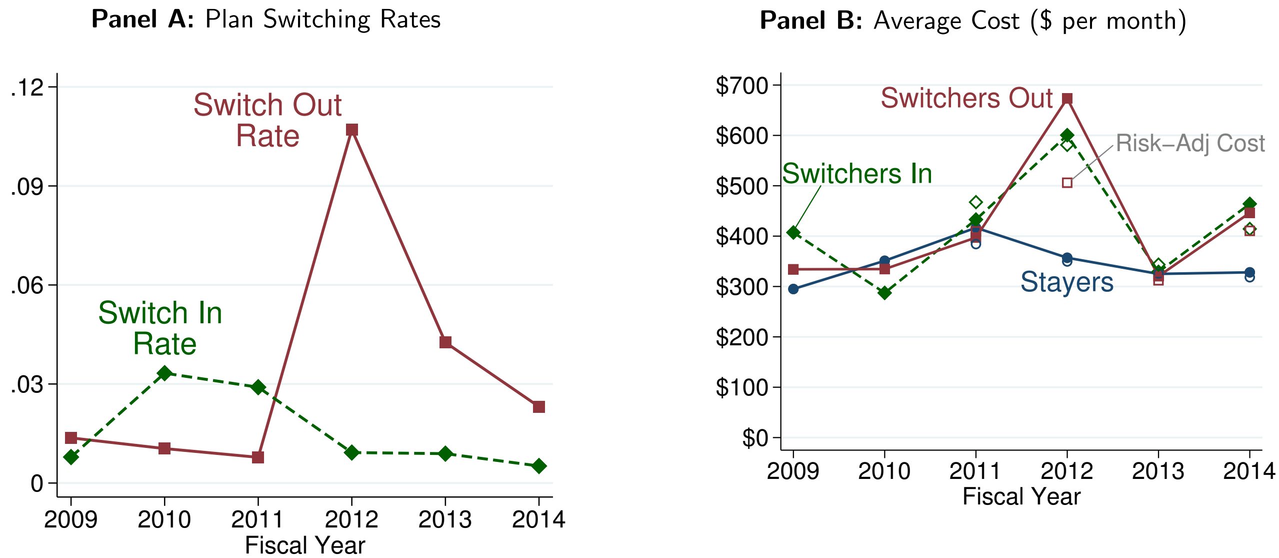

NOTE: These figures show switching and selection patterns for *zero-premium* (below-poverty) Network Health over time and especially around its 2012 network narrowing. The graphs are exactly analogous to Figure 2 in the main text but with the sample limited to below-poverty enrollees who do not pay premiums.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

Figure A.5: Plan Switching Out Rates for Network Health: Zero-Premium Enrollees

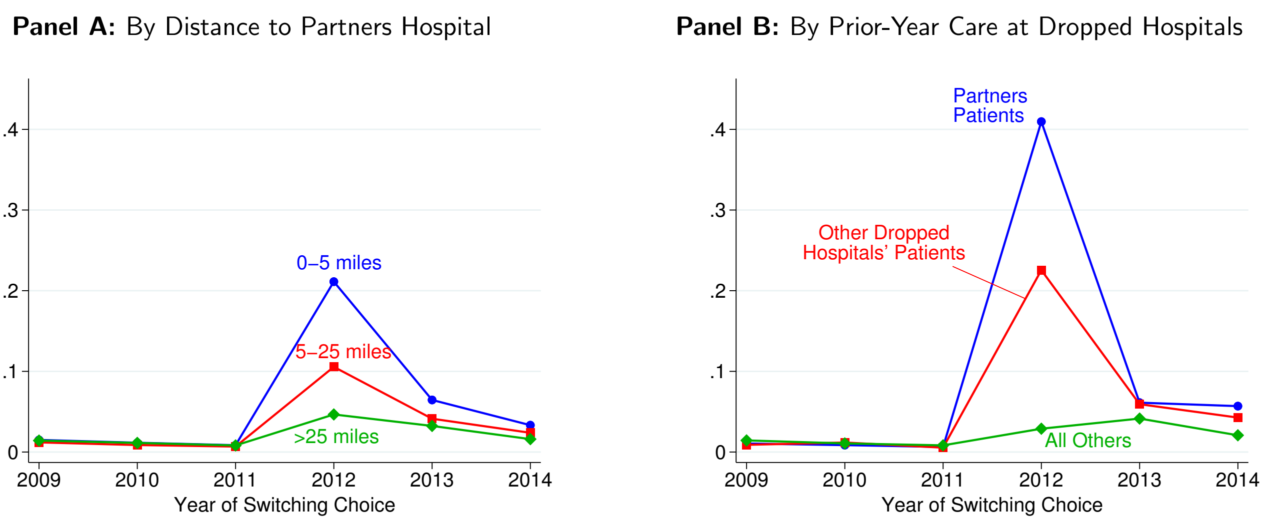

NOTE: These figures show switching out patterns for *zero-premium* Network Health enrollees around its 2012 dropping of Partners and several other hospitals. They are exactly analogous to Figure 2 in the main text but with the sample limited to below-poverty enrollees who do not pay premiums. See the caption for Figure 2 for additional information.

(2) 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.

Appendix Figure A.6 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 2 in the main paper. 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 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 about 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 A.7 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.

Figure A.6: Network Health's New Enrollee Market Share around 2012 Change

Panel A: By Distance to Partners

Panel B: By Prior Provider Use

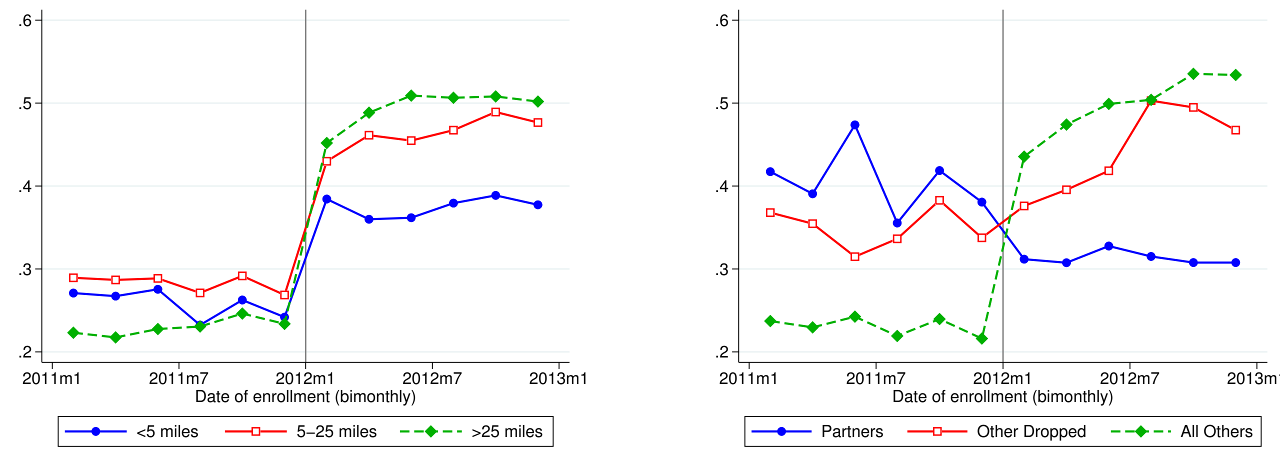

NOTE: These figures show evidence of changes in new enrollees' demand for Network Health in 2012 that are correlated with valuation for the Partners and other dropped hospitals. Each point on the figures is the market share who choose Network Health among above-poverty new enrollees joining the exchange 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.

Figure A.7: Changing Risk Selection for Network Health among New Enrollees

Panel A: Average Risk Score (all new enrollees)

Panel B: Prior-Spell Costs ($/month, re-enr. only)

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.

(3) 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-A.6 above for CeltiCare.

The results are shown in Appendix Figures A.8-A.10. 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.

Figure A.8: Plan Switching and Selection for CeltiCare (Drops Partners in 2014)

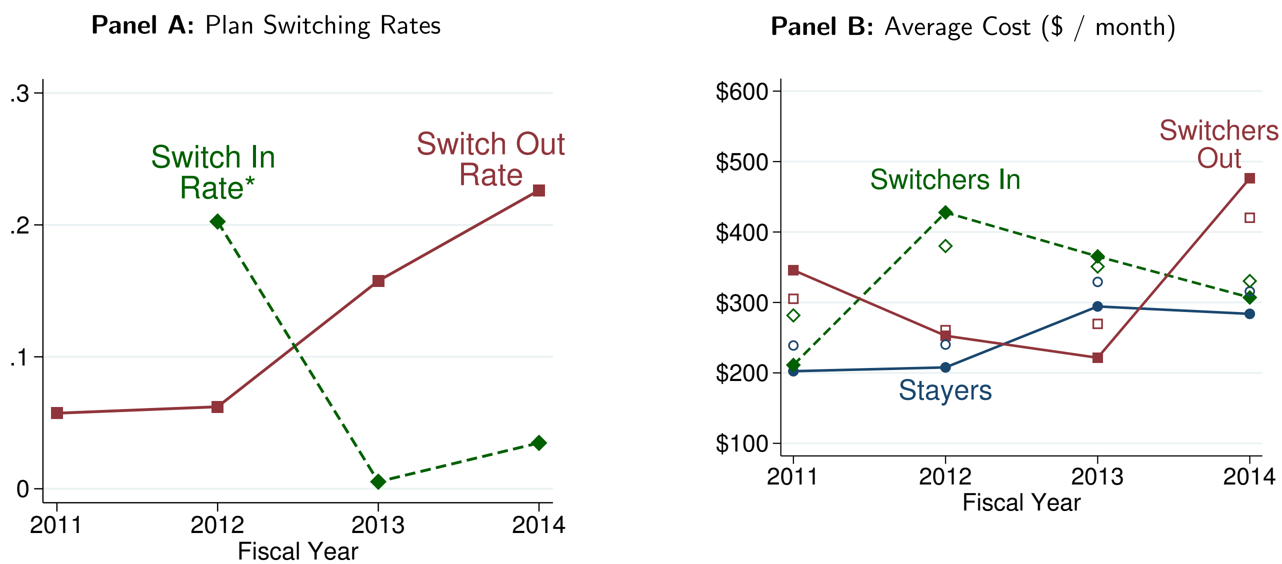

* Panel A excludes the 2011 switching in rate for CeltiCare to avoid blowing up the y-scale.
NOTE: These figures show switching rates for CeltiCare (Panel A) and average prior-year costs for CeltiCare enrollees (Panel B, in $ per month) in each year's open enrollment. CeltiCare drops the Partners Healthcare system from its network in 2014. These plots are analogous to Figure 3 in the main text and Appendix Figure A.12, which show switching and selection for Network Health. See the notes to those figures for additional description. The current figure shows that similar adverse selection patterns occur for CeltiCare when it excludes Partners from network.

Figure A.9: Switching Out Rates for CeltiCare (Drops Partners in 2014), by Enrollee Characteristics

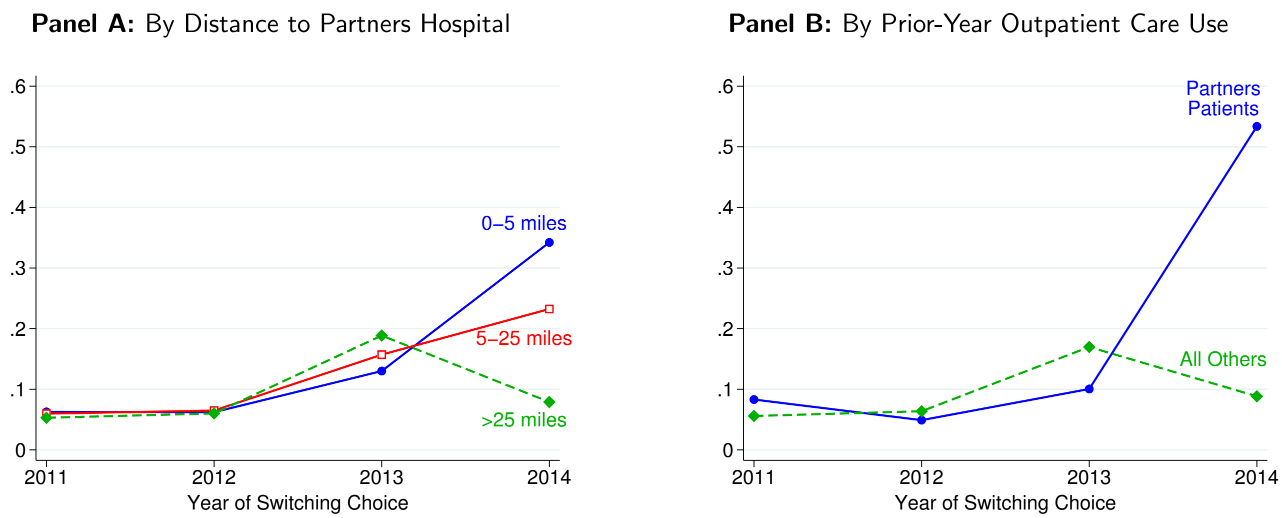

NOTE: These figures show switching out rates for CeltiCare enrollees by variables likely to correlate with demand for Partners, which is dropped from the plan's network in 2014. Panel A shows switching rates by enrollee distance to the nearest Partners hospital; Panel B shows switching rates by prior-year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

Figure A.10: CeltiCare's Market Share among New Enrollees (Drops Partners in 2014)

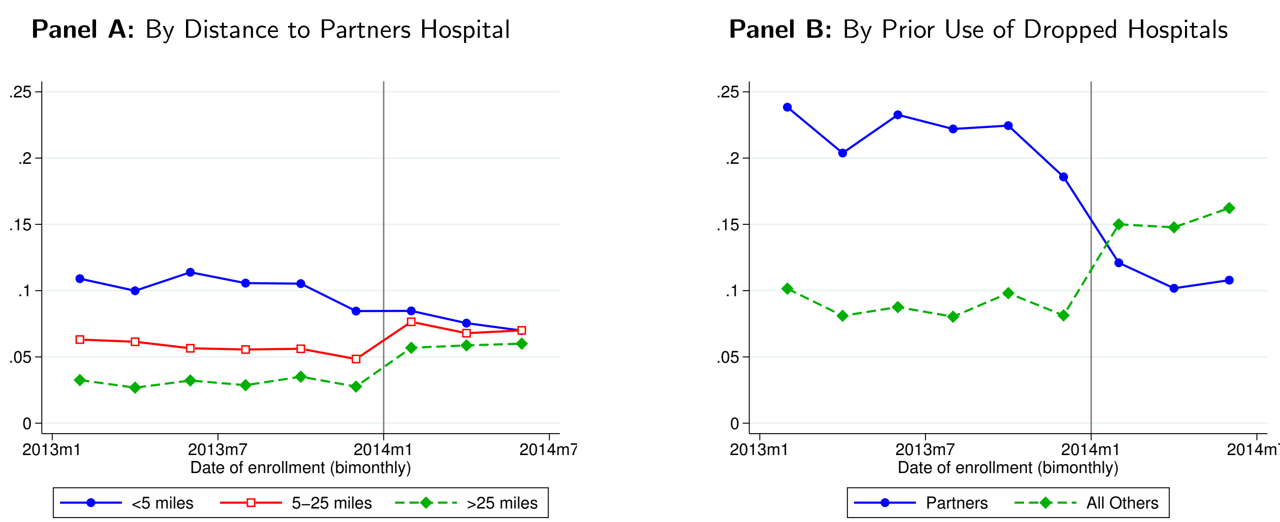

NOTE: These figures show evidence of changes in new enrollees' demand for CeltiCare in 2014 (when it drops Partners from network) that are correlated with valuation for Partners providers. (The plots are analogous to Figure A.6 in the main text, which studies Network Health's network change in 2012.) 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 A.11: Monthly Rate of Exiting the Exchange, Network Health Enrollees

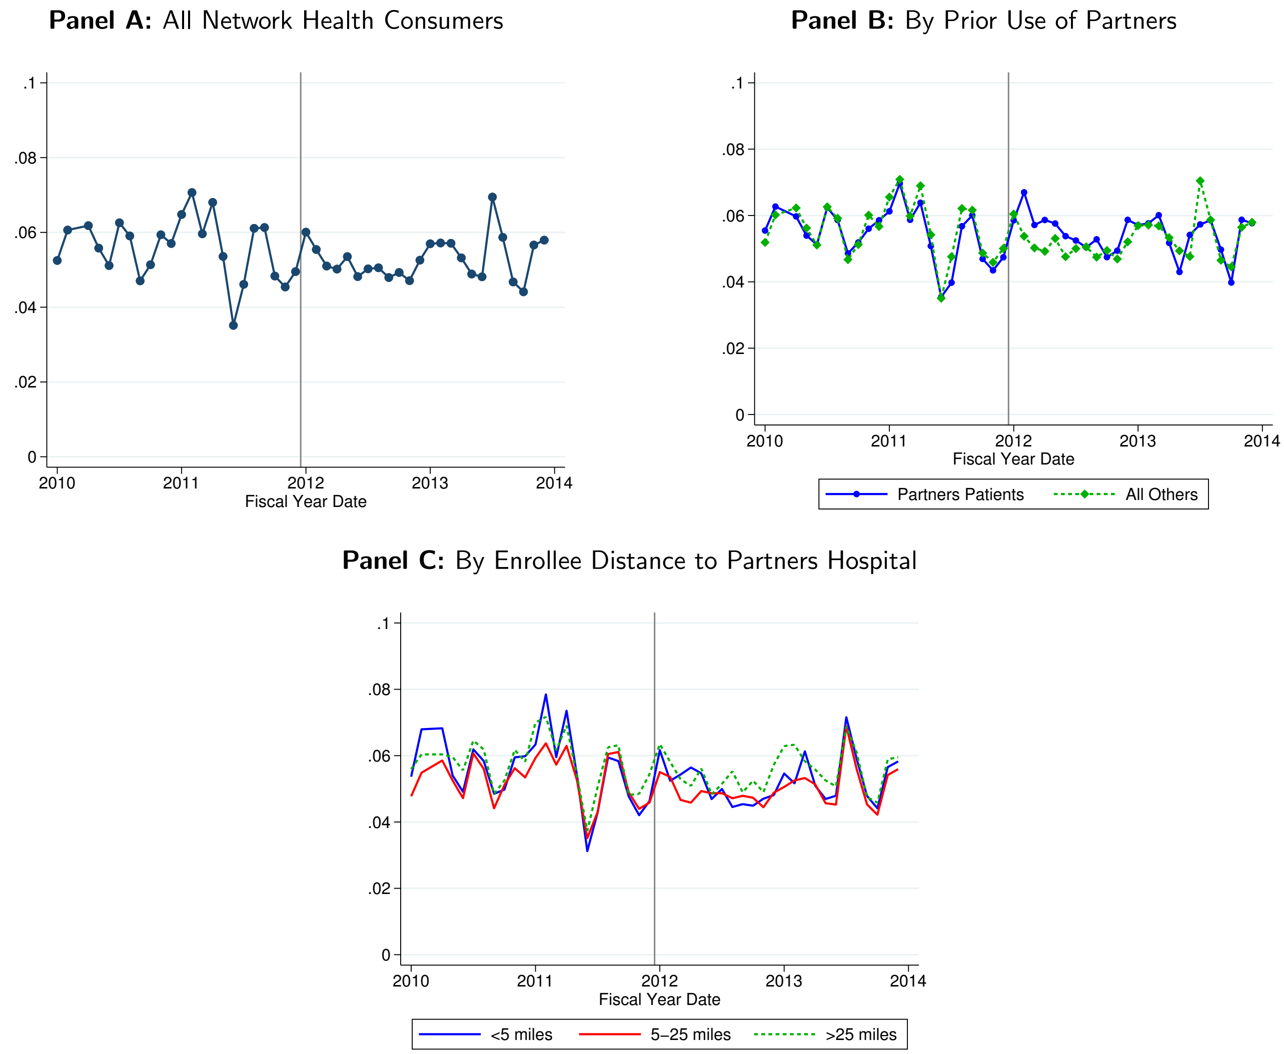

NOTE: The figure provides evidence on a key assumption in the plan choice model: that Network Health's network narrowing in 2012 does not affect whether consumers participate in the exchange (no "extensive margin" response). The figure plots the share of Network Health's existing enrollees who exit the 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).

C.2 Additional Analyses on Reduced Form Switching and Selection Patterns

This appendix shows additional facts about plan switching and selection into and out of Network Health and runs robustness checks on the excess switching rate logits shown in Section 4.1.

1. Switching Rates In and Out of Network Health Figure A.12 shows switching rates for Network Health in each year from 2009-2014. 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 switchers *into* the plan divided by the same denominator, which allows for comparing the two figures in levels. At the start of 2012 when its narrower network (and lower price) took effect, the plan experienced a spike in switching – to 11.3% 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.[64] 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.

2. Breakdown of Costs of Switchers and Stayers by Group Appendix Table A.4 shows evidence that the groups most likely to switch out of Network Health in 2012 also have high costs, implying adverse selection. 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.

3. Robustness Check on Logit Regressions for Switching Patterns Figure A.13 shows a robustness check on Figure 4 in the body text. It shows estimates from a multivariate version of the logit regression in equation (5), with distance, observed sickness (quantile of the CommCare risk score), and unobserved sickness (ratio of HCC risk score to the CommCare risk score) all included as covariates in the same specification. The results are estimates of the odds ratio for excess switching in 2012 ($= \exp(\beta_g)$ in equation (5)). The results confirm that distance, observed risk, and unobserved risk all separately predict plan switching in 2012 in a multivariate specification.

Figure A.14 shows another robustness check on these logit regressions. It replicates the top three panels of Figure 4 in the body text, separately for prior-year Partners patients (red triangles) and people who were not patients of a dropped hospital (blue circles). Distance, sickness, and unobserved sickness continue to predict plan switching in 2012 within each subgroup, though the sickness gradient is stronger for the Partners patients and the distance gradient is somewhat stronger for non-patients.

[64]Switching out rates were also high in 2009, reflecting unusually large increases in Network Health's enrollee premiums from 2008-09.

Figure A.12: Plan Switching Rates In and Out of Network Health (around 2012 network change)

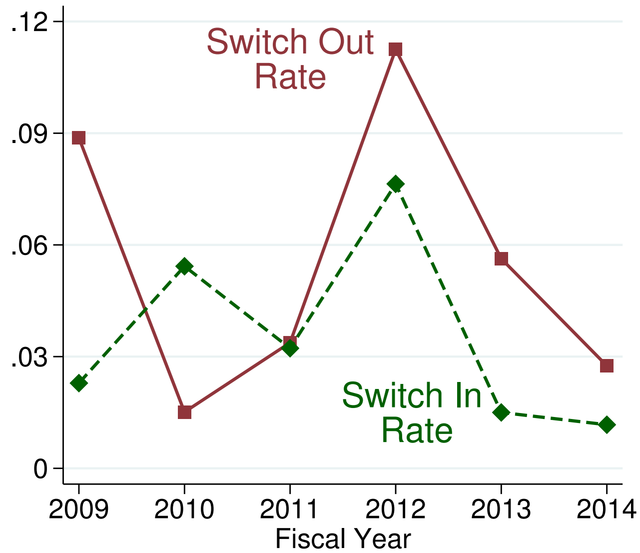

NOTE: The figure shows switching patterns for Network Health over time and especially around its 2012 network narrowing. It plots the rate of switching in and out of Network Health at each year's open enrollment. These rates are defined as the number of switchers in/out divided by the same denominator – the number of continuous market enrollees in Network Health at the end of the prior year – so their levels are comparable.

Table A.4: 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 (1)	Risk Adj. Cost (2)	Share of Enrollees (3)	Switching Rate (4)	Share of Switchers (5)	2011 (6)	2012 (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.

Figure A.13: Excess Switching Out Rates in 2012: Multivariate Logit Estimates

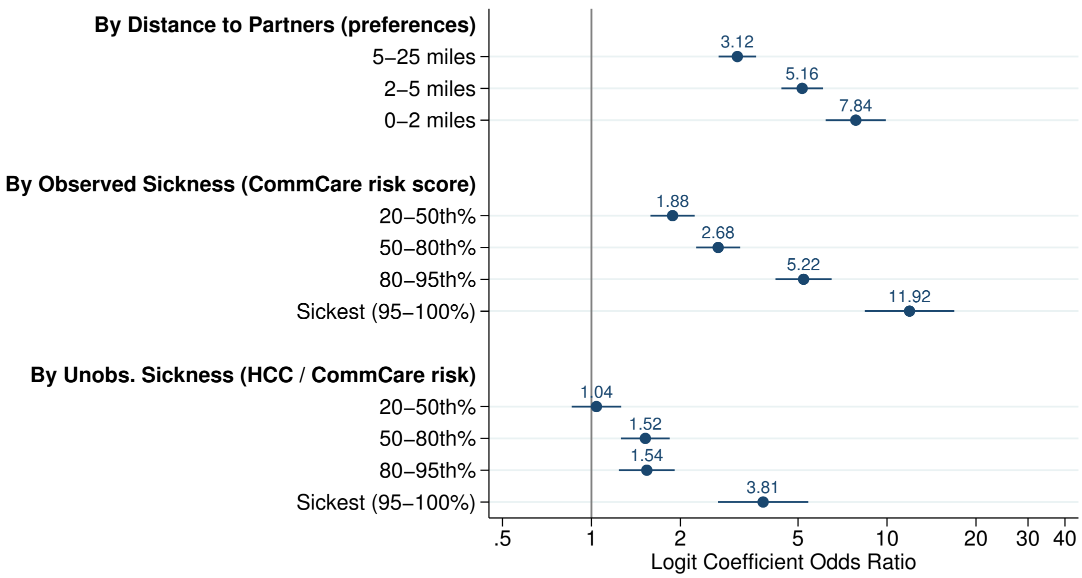

NOTE: The figure shows odds ratios corresponding to $\exp(\beta_g)$ from estimates of switching multivariate logit regression specification (5). The results come from a single logit regression with distance, observed sickness, and unobserved sickness as covariates. Distance is defined as enrollee distance to the nearest Partners hospital (with an omitted group of 25+ miles). Observed sickness is defined as quantiles of the (prior-year) CommCare risk score (with 0-20th% as the omitted group), which is the measure used for actual risk adjustment. Unobserved sickness is defined as the ratio of of the HCC risk score to CommCare's risk score, both measures for the prior year.

Figure A.14: Excess Switching Out Rates in 2012: Separately by Past Patient Status

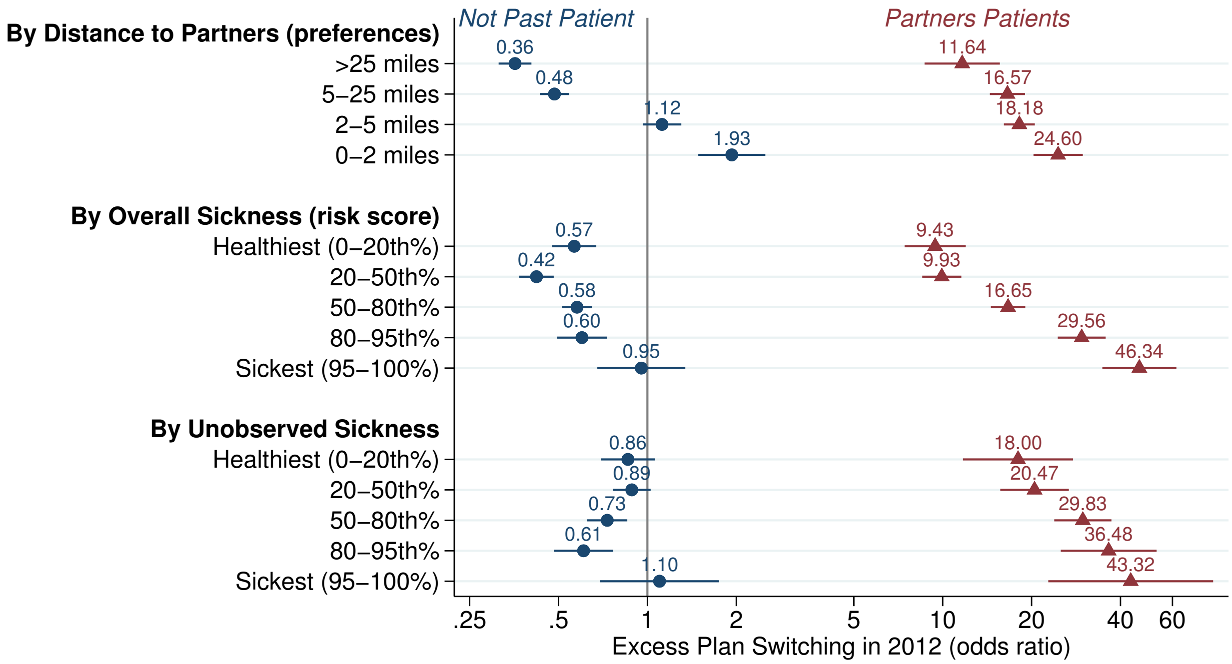

NOTE: The figure shows odds ratios corresponding to $\exp(\beta_g)$ from estimates of switching multivariate logit regression specification (5). The results come from a single logit regression with distance, observed sickness, and unobserved sickness as covariates. Distance is defined as enrollee distance to the nearest Partners hospital (with an omitted group of 25+ miles). Observed sickness is defined as quantiles of the (prior-year) CommCare risk score (with 0-20th% as the omitted group), which is the measure used for actual risk adjustment. Unobserved sickness is defined as the ratio of of the HCC risk score to CommCare's risk score, both measures for the prior year.

D Appendix: Understanding Demand for Star Providers

D.1 Decomposition of Role of Sickness vs. Preferences in Demand

To quantify the role of sickness versus preference measures in explaining demand for the star Partners hospitals, I implement a decomposition method suggested by Shorrocks (2013); see also Shorrocks (1982). The method, which is also known as a "Shapley-Shorrocks decomposition," quantifies the role of covariates in explaining variation in an outcome variable.[65] This role is quantified by the marginal contribution of a covariate (or group of covariates) to the $R^2$of a regression – i.e., how much the R^2 increases when a covariate is added. To account for complementarity among covariates (which means that the ordering in which covariates enter matters), it calculates the Shapley value of this contribution – essentially averaging over the marginal contribution to R^2 for every possible covariate ordering. I implement the method using the add-on Stata command "shapley2".

I implement the decomposition for two metrics of demand (Y_i) for Partners: (1) switching plans in 2012, and (2) being a Partners patient in 2011. I restrict the sample to the 2012 current enrollee sample enrolled in Network Health at the end of 2011. I run logit regressions of the form:

$$Y_i = logit\left(\alpha + X_i^{Dist}\beta_1 + X_i^{Sickness}\beta_2\ [+X_i^{ProvRelat}\beta_3]\right)$$

where X_i^{Dist} is a vector of covariates for distance to the nearest Partners hospital (10 deciles up to 35 miles away, plus a dummy for 35+ miles) and to the nearest other dropped hospital (similar variables); $X_i^{Sickness}$ is a vector of sickness covariates, including "observed" and "unobserved" risk; and $X_i^{ProvRelat}$ are dummies for being a patient of Partners and of another dropped hospital during 2011 (only included when Y_i = switching plans). Observed risk covariates include age groups and deciles of the CommCare risk score, plus an extra category for the top 5%. I consider two versions of unobserved risk. A simpler version includes quantiles of the HCC risk score (deciles + top 5% dummy) and dummies for nine chronic illnesses. A richer version includes these variables plus variables for prior-year (2011) utilization of care (e.g., quantity of care, number of office visits, any hospitalization) and subsequent-year (2012) HCC risk score quantiles and diagnosis variables, which can capture the role of future health shocks.[66] The bottom of Table A.5 reports the number of variables for each group of covariates.

Table A.5 reports results of the decomposition for four covariate specifications: (1) distance + observed risk only, (2) adding the simpler unobserved risk covariates, (3) adding the richer unobserved risk covariates, and (4) adding provider relationships. The first panel shows results for the demand measure (Y_i) of switching out of Network Health in 2012; the second panel shows results for being a Partners patient in 2011. In each panel, the top row lists the overall explained variation (McFadden's pseudo-R^2) and the contribution of each set of covariates to this R^2 (these by construction add up to

[65] The method is sometimes used to quantify the contribution of factors to explaining distributional inequality (e.g., in income or wealth). It is distinct from the better known Oaxaca-Blinder decomposition, which decomposes the role of factors in explaining inequality *between two groups* (e.g., the black-white income gap).

[66] Although subsequent-year variables are potentially endogenous to the switching choice, this very rich specification allows me to capture any future health shocks that emerge during 2012 and that agents might have known when making switching decisions.

the total).[67] As noted above, the covariate contribution represents the average marginal increase in the psuedo-R^2 when this group of covariates is added to the specification (i.e., the Shapley value of their contribution).

The results in Table A.5 suggest that while preferences and sickness both matter, preferences are quantitatively more important in explaining demand variation. Even in the richest specification for sickness (column 3, which includes 64 sickness covariates), distance accounts for 56% of the explained variation in switching plans and 69% of the explained variation in being a Partners patient, with sickness variables accounting for the remainder. There is also substantial unobserved variation, as indicated by the pseudo-R^2 of 0.147-0.285. Although this unexplained variation may reflect either unobserved preferences or sickness, unobserved preferences are likely more important. Distance is just one driver of preferences, while sickness is relatively well measured in claims data. Moreover, column 4 shows that adding provider relationship dummies (just two variables) more than doubles the R^2 to 0.336, and these dummies account for more variation than all of the distance and sickness variables combined.

D.2 Role of State Dependence vs. Heterogeneity

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[68] – 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 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

[67] I use the psuedo-R^2 because this is a logit regression, but I have found that results are nearly identical if I instead run a linear probability model and use the traditional R^2.

[68] 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 A.3) and a worse reputation (as indicated in the plan demand estimates in Table A.11). Interestingly, switching rates for below- and above-poverty enrollees were quite similar.

Table A.5: Role of Sickness vs. Preferences in Explaining Demand for Star Hospitals

	Observed Risk Only		Add Unobs. Risk		Additional Risk Covars.		With Provider Relationships	
	(1)		(2)		(3)		(4)	
Demand Measure #1: Switching Out of Network Health in 2012								
Explained Variation (McFadden's Pseudo-R^2)	**0.106**		**0.130**		**0.147**		**0.336**	
Contribution to Pseudo-R^2								
Distance to dropped hospitals (preference)	0.083	*[79%]*	0.083	*[64%]*	0.083	*[56%]*	0.054	*[16%]*
Sickness: Observed (in risk adjustment)	0.022	*[21%]*	0.013	*[10%]*	0.012	*[8%]*	0.011	*[3%]*
Unobserved (not in risk adj.)	---		0.035	*[27%]*	0.052	*[35%]*	0.038	*[11%]*
Patient of dropped hospitals	---		---		---		0.234	*[69%]*
Demand Measure #2: Being a Partners Patient in 2011								
Explained Variation (McFadden's Pseudo-R^2)	**0.204**		**0.276**		**0.285**		---	
Contribution to Pseudo-R^2								
Distance to dropped hospitals (preference)	0.192	*[94%]*	0.196	*[71%]*	0.197	*[69%]*		
Sickness: Observed (in risk adjustment)	0.012	*[6%]*	0.007	*[3%]*	0.007	*[2%]*		
Unobserved (not in risk adj.)	---		0.073	*[26%]*	0.080	*[28%]*		
Covariates Included								
Distance to Partners, other dropped hosp. (n = 20)	X		X		X		X	
Prior-Year Patient of Dropped Hospitals (n = 2)							X	
Sickness covariates								
Age groups (n = 9)	X		X		X		X	
CommCare risk score bins (n = 10)	X		X		X		X	
HCC risk score bins (n = 10)			X		X		X	
Diagnoses dummies (n = 9)			X		X		X	
Prior-Year Utilization variables (n = 5)					X		X	
Subsequent-year risk score & diagnoses (n = 21)					X		X	
Number of Observations	41,917		41,917		41,917		41,917	

NOTE: The table reports results of the Shorrocks decomposition of the contribution of distance and sickness covariates to explained variation (the pseudo-R^2) in two demand outcomes: (1) switching out of Network Health in 2012 when it drops Partners (top panel), and (2) being a Partners patient in 2011 (middle panel). See the appendix text for a detailed description of the method for this decomposition. The sample is restricted to current enrollees in Network Health as of the end of 2011, just as in the reduced form analysis in the paper.

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 A.6 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 A.6 indicate support for both 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.

Table A.6: Heterogeneity in Likelihood to Switch Out after 2012 Network Narrowing

Variable	Outcome: Switch Out of Network Health — Unconditional (1) Odds Ratio	(S.E.)	Outcome: Switch Out of Network Health — 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.

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 A.15 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).[69] 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 A.7 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, 43-45% 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 A.15 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

[69] 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 A.15: Switching Rate Out of Network Health, by Recency of Last Provider Visit

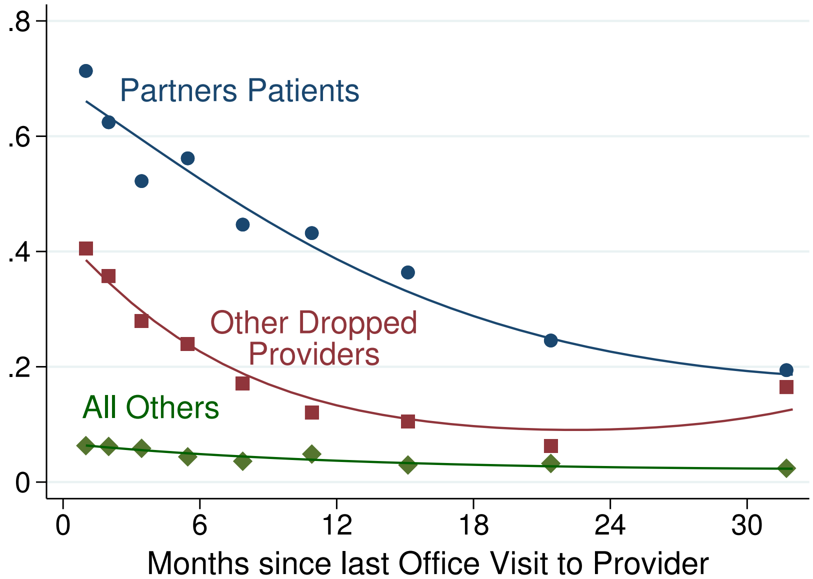

NOTE: The plot shows how plan switching rates out of Network Health in 2012 relate to the recency of a physician office visit with the indicated provider. Individuals are categorized into Partners patients (blue circles), patients of another dropped hospital (red squares), and all other patients (green 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.

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 A.7). 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.

Table A.7: 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) (1) Prob.	(S.E.)	Medical Risk Controls (2) Prob.	(S.E.)	Medical Risk and Distance Controls (3) 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 A.15 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 A.6. Separate regressions are run for each patient group.

E Appendix: Cost Decomposition Details and Analyses

This appendix describes additional details of the method for decomposing medical spending, as summarized in Section 5.1, and also presents additional analyses related to the findings in Section 5.

E.1 Cost Decomposition Method Details

As discussed in Section 5.1, I decompose costs into prices vs. quantities, and quantities into risk-predictable quantity and a residual. The method involves four key steps:

1. Defining the unit of medical services (s)
2. Estimating the "quantity" of 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
4. Estimating Risk-Predictable Quantity

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

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.[70] I also exclude outpatient emergency department care to avoid double-counting, since these are included in the inpatient costs 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

[70] 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.

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.

Let A_{it}^{OP} be the set of outpatient services used by person i in year t, and let a_{it} index each instance of utilization. With these definitions (and following Section 5.1), total quantity of outpatient care for an enrollee equals

$$Q_{i,t}^{OP} = \sum_{a_{it} \in A_{it}^{OP}} Q_{s(a_{it})} = \sum_{a_{it} \in A_{it}^{OP}} \overline{Paid}_{s(a_{it})}. \tag{11}$$

Average price equals the residual factor explaining costs, which is also a (quantity-weighted) average of prices across all services used by the individual:

$$P_{i,t}^{OP} \equiv \frac{C_{i,t}^{OP}}{Q_{i,t}^{OP}} = \sum_{a_{it} \in A_{it}^{OP}} \left(\frac{Q_{s(a_{it})}}{Q_{it}^{OP}} \right) \cdot P_{a_{it},s}. \tag{12}$$

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.[71] 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:[72]

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

Using estimates of (13), I define the quantity unit as the component of payment arising from DRG/diagnosis,

[71] 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 (AHRQ). 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.

[72] 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.

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 definitions of price and quantity, I apply the same idea as in equations (11) and (12) for outpatient care to define inpatient quantity for i in year t as $Q_{i,t}^{IP} \equiv \sum_{a\in Admit(i,t)} \widetilde{Q}_{a,i,t}$, and price as $P_{i,t}^{IP} \equiv C_{i,t}^{IP}/Q_{i,t}^{IP}$.

Combined Inpatient and Outpatient Costs

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:

$$Q_{i,t}^{Tot} \equiv Q_{i,t}^{IP} + Q_{i,t}^{OP} \tag{16}$$

Price is defined as the remaining factor needed to account for costs (which as noted above equals a weighted average of service-level prices):

$$P_{i,t}^{Tot} = \frac{C_{i,t}^{IP} + C_{i,t}^{OP}}{Q_{i,t}^{Tot}} \tag{17}$$

Estimating Risk-Predictable Quantity

After pulling out quantity, I project it (separately for outpatient and inpatient care) onto medical risk observables (Z_{it}) to estimate "risk-predictable quantity." To deal with the combination of zeros and skewed distribution of Q_{it}, I estimate a two-part model, with a logit for the probability of positive quantity and log-linear regression for quantity conditional on positive. Specifically, the two parts are: (1) the logit model: $Pr\left(Q_{it} > 0\right) = Logit\left(Z_{it}\theta_1\right)$, and (2) the log-linear model: $\log Q_{it}|Q_{it} > 0 = Z_{it}\theta_2 + \varepsilon_{it}$. These models are estimated using the Stata command "twopm". The command uses the estimates to output predicted quantity as:

$$\hat{Q}_{it}^{risk} = E\left[Q_{it}|Z_{it}\right] = Logit\left(Z_{it}\hat{\theta}_1\right)\cdot\exp\left(Z_{it}\hat{\theta}_2\right)\cdot E\left(e^{\varepsilon}\right)$$

where $Logit\left(.\right) = \frac{\exp(.)}{1+\exp(.)}$ and the $E\left(e^{\varepsilon}\right)$ is the "Duan smearing" correction so that the mean of $\hat{Q}_{it}$ more closely matches Q_{it}, a method that works better than using the standard log-normal factor $\exp\left(\sigma_{\varepsilon}^2/2\right)$.[73]

[73] See the documentation for Stata's "twopm" command for additional details.

I do this projection first using only variables included in the exchange's (retrospective) risk adjustment, including age and a flexible 11-part spline for the CommCare risk score. This generates what I call "observed risk": $\hat{Q}_{it}^{risk,obs} = f\left(Z_{it}^{obs}; \hat{\theta}\right)$. I then do the decomposition for these variables plus a broader set of risk variables from the claims, including concurrent diagnoses and a spline of the concurrent HCC risk score. This generates my overall measure of risk-predictable quantity: $\hat{Q}_{it}^{risk} = f\left(Z_{it}^{obs}, Z_{it}^{other}; \hat{\theta}\right)$. I then define "residual quantity" as the remaining factor explaining observed quantity: $\hat{Q}_{it}^{resid} \equiv Q_{it}/\hat{Q}_{it}^{risk}$.

Summary of Decomposition

Putting everything together, individual-level costs equal the product of three factors: $C_{it} = \hat{Q}_{it}^{risk} \cdot \hat{Q}_{it}^{resid} \cdot P_{it}$. This relationship also holds at a group level for (appropriately weighted) averages:

$$\overline{C}_{g,t} = \overline{Q}_{g,t}^{risk} \times \overline{Q}_{g,t}^{resid} \times \overline{P}_{g,t} \tag{18}$$

where $\overline{P}_{g,t}$ is average prices weighted by enrollee quantity (Q_{it}), and $\overline{Q}_{g,t}^{resid}$ is the average residual weighted by risk-predicted quantities ($\hat{Q}_{it}^{risk}$). This equation lets me decompose the share of group cost differences (e.g., stayers vs. switchers in 2012) that are driven by (1) risk-predictable quantity, (2) residual quantity, and (3) provider prices. Its multiplicative form suggests decomposing log differences for each factor, which are additive:

$$\Delta \log\left(\overline{C}\right) = \Delta \log\left(\overline{Q}^{risk}\right) + \Delta \log\left(\overline{Q}^{resid}\right) + \Delta \log\left(\overline{P}\right)$$

This allows me to quantify the share of log cost differences explained by these three factors, as shown in Table (2).

E.2 Summary Statistics on Price-Quantity Estimates

Appendix Table A.8 shows summary statistics from the decomposition. Panel A shows statistics about the mean and standard deviation of medical costs and the quantity and price decomposition estimates. In addition to quantity in dollars per month, I show statistics for quantity relative to the sample mean, to make the units more comparable to the price variable. Panel B shows the relationship of quantity and price to the HCC 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 around the network change. 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 \$228.2 per month (which, is 61% of overall average costs of \$375). Its standard deviation of \$780 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 (coefficient of variation = 0.33). 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 HCC enrollee risk scores, using simple regressions of quantity/price on risk score and a constant. (The HCC risk score is a concurrent measure used by the ACA and capture more information about risk than the retrospective CommCare risk score, especially for new enrollees.) 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 (scaled relative to the sample mean) – with a regression coefficient of 0.408 (s.e. = 0.007) – they hardly predict price variation at all (coeff. = -0.0004, s.e. = 0.0001). Similarly, the R^2 is about 26% for quantity versus <0.1% for price. The pattern is similar for outpatient quantity. Risk score is slightly better at predicting inpatient prices, with a coefficient of 0.002 and R^2 of 1.3%, but these are still an order of magnitude smaller than the analogs for inpatient quantity.

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

Table A.8: Price vs. Quantity Medical Cost Decomposition

Variable	Statistic	Total Costs (1)	Outpatient Costs (2)	Inpatient Costs (3)
A. Cost Decomposition Summary				
Costs in Decomp.	Mean	$228.2	$163.6	$64.7
($ per month)	*[S.D.]*	*[$779.5]*	*[$388.7]*	*[$609.2]*
Quantity of Care	Mean	$228.6	$165.9	$62.7
($ per month)	*[S.D.]*	*[$720.9]*	*[$395.8]*	*[$536.2]*
Quantity (relative	Mean	1.00	1.00	1.00
to mean)	*[S.D.]*	*[3.15]*	*[8.55]*	*[2.39]*
Price Factor	Mean	1.02	1.02	1.00
	[S.D.]	*[0.34]*	*[0.34]*	*[0.26]*
B. Regression of Quantity/Price on Risk Score				
Quantity (relative	Regr. Coeff	0.408	0.912	0.217
to mean)	(s.e.)	(0.007)	(0.024)	(0.004)
	[R^2]	*[26.0%]*	*[17.7%]*	*[12.9%]*
Price Factor	Regr. Coeff	-0.0004	-0.0010	0.0022
	(s.e.)	(0.0001)	(0.0001)	(0.0001)
	[R^2]	*[0.00%]*	*[0.02%]*	*[1.27%]*

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), for quantity (in $ per month) and quantity relative to the sample mean, and for price (see text for its definition). Panel B shows estimates of regressions of quantity/price (y-variable) on an enrollee's HCC risk score. For both panels, the columns show results separately for (1) total costs in the decomposition (outpatient + inpatient costs), (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 surrounding the key network change.

E.3 Switchers vs. Stayers Costs: Additional Analyses

Table 2 in the body text quantifies the contribution to switcher-stayer cost differences of overall quantity, risk-predictable quantity, residual quantity, and provider prices. These differences shed light on the role of the two cost dimensions, medical risk and provider costs/choices, to cost differences driving adverse selection. This appendix discusses additional analyses that illustrate how switchers and stayers differ in different components of the cost decomposition.

Descriptive Plots of Stayers vs. Switchers Cost Differences Figures A.16 and A.17 show descriptive plots on the distribution of components of the cost decomposition for stayers vs. switchers out of Network Health in 2012. In most panels, the left figure shows the overall measure's distribution for stayers (red) vs. switchers (blue), while the right panel shows a bin scatter of the mean by decile of enrollee HCC risk score, to illustrate the risk-conditional distribution. Each panel shows a different variable relevant to the cost decomposition: total medical spending included in the decomposition (panel A), total quantity of care (panel B), risk-predictable quantity of care (panel C) using all risk variables, inpatient prices (panel D), outpatient prices (panel E), and share of utilization that occurs at Partners providers (panel F). (In panel F, I do not show the overall distribution, which is bimodal and hard to see, but instead show risk bin scatters separately for inpatient and outpatient costs.)

The figures indicate that switchers are higher cost on nearly all metrics. Switchers' overall higher costs (panel A) are evident in the raw distributions, with stayers having a much higher density peak at low spending levels. The left figure of panel A shows that the differences are consistent across the risk distribution and close to constant in percentage terms (note the log scale of the axes). In particular, switchers' costs are higher than stayers regardless of whether they are healthy or sick. Panel B shows that a similar pattern holds for overall quantity of care, which drives the majority of cost differences. Panel C shows that risk-predictable quantity captures much of the differences, but there is still a gap as indicated by the higher residual quantity for switchers (see Table 2). Some of this residual quantity may reflect effects of provider treatment intensity.

Figure A.17 show that prices and provider use is also relevant. Panel A shows that switchers have almost 30% higher inpatient care prices, which is entirely driven by their much higher likelihood to choose Partners hospitals (panel F1). Panel B shows, by contrast, that switchers and stayers have similar outpatient care prices. Even though switchers are much more likely to choose Partners for outpatient care (panel F2), I estimate that Partners outpatient prices are not high, leading to the result of similar outpatient prices.

Analysis of Residual Quantity The estimates in Table 2 indicate that residual quantity (not predictable by medical risk variables) accounts for a meaningful share of the higher costs of switchers out of Network Health relative to stayers. This residual quantity is challenging to interpret because it might reflect either further unobserved medical risk (the standard cost channel) or provider effects on treatment intensity (the second channel). To provide suggestive evidence on this issue, Table A.9 analyzes whether residual quantity is associated with provider use, proxied by being a patient of Partners or another dropped hospital. Because patients may sort on unobserved medical risk,

Table A.9: Analysis of Residual Quantity and Partners Use

	Outcome: Residual Quantity			
	Outpatient Care		**Inpatient Care**	
	OLS	Distance IV	OLS	Distance IV
	(1)	(2)	(3)	(4)
Partners Patient	0.292**	0.158**	0.299**	0.147
	(0.020)	(0.046)	(0.066)	(0.149)
Other Dropped Hospital Patient	0.221**	0.034	0.134	0.346
	(0.023)	(0.069)	(0.070)	(0.213)
Constant	0.770**	0.823**	0.758**	0.779**
	(0.007)	(0.019)	(0.030)	(0.061)
First-Stage F-Stats.				
Partners Patient	---	219.3	---	41.0
Other Dropped Patient	---	107.0	---	26.6
Num. Obs.	41,917	41,917	41,917	41,917

NOTE: The table shows estimates of a regression of residual quantity on dummies for being a Partners patient or patient of another hospital dropped by Network Health in 2012. The sample is all current Network Health enrollees at the end of 2011 (same as for Table 2 in the main text), and the residual quantity measure is for 2011 and is defined by the price-quantity decomposition in Section 5.1. Columns (1)-(2) show estimates for outpatient costs, while columns (3)-(4) show inpatient costs. Columns (1) and (3) show OLS estimates, while columns (2) and (4) instrument for patient status using distance to the relevant hospitals. If distance is orthogonal to unobserved medical risk, these IV estimates reflect causal provider impacts on quantity of care.

columns (2) and (4) instrument for patient status using distance to the relevant provider. If distance is orthogonal to unobserved medical risk, then these IV estimates represent causal provider impacts on quantity.

The IV estimates suggest that Partners patients have about 15% points higher residual quantity of outpatient care (significant at the 1% level) and a non-significant 15% higher residual quantity for inpatient care (where the estimates are much noisier). These estimates are about 20% of the average for other enrollees (captured by the regression constant). The IV estimates are about half of the OLS estimates, suggesting that high residual quantity reflects a mixture of unobserved medical risk and Partners provider impacts on quantity.

Figure A.16: Switcher vs. Stayer Cost Decomposition: Distributions and Bin Scatters by Risk Score

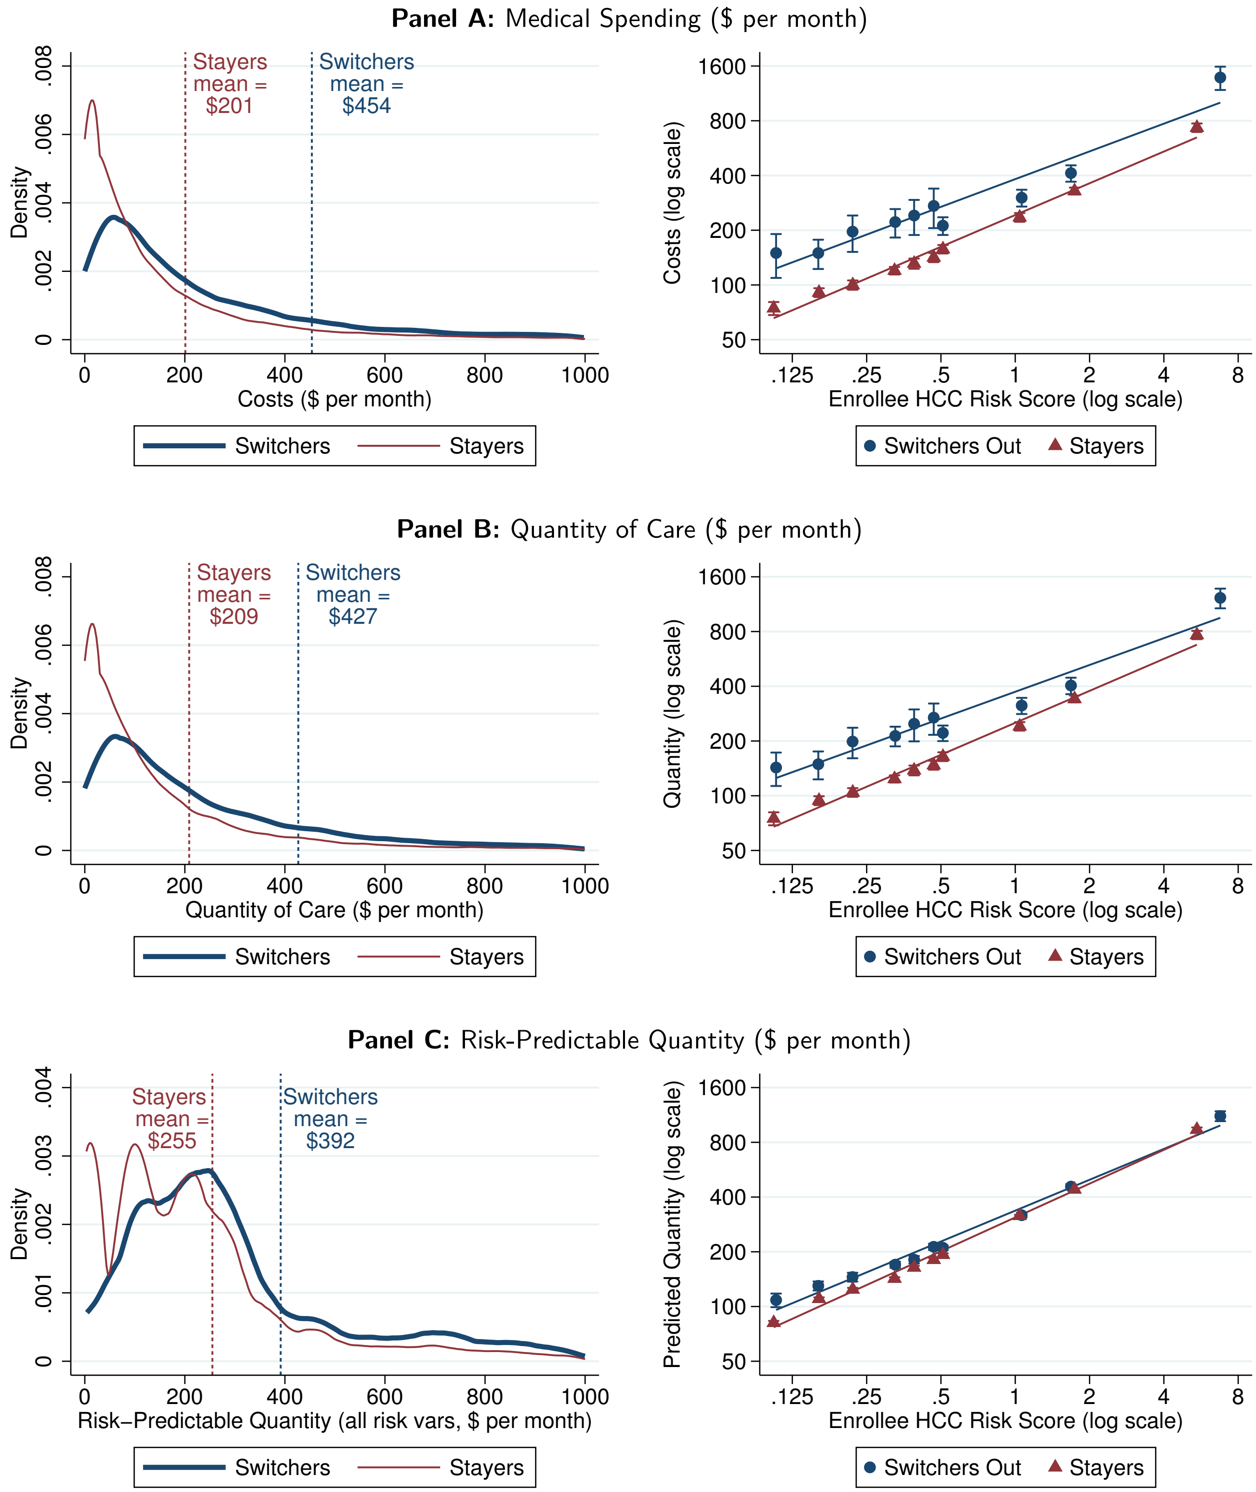

NOTE: The figures show the distribution and risk score-conditional distributions of cost components for switchers vs. stayers in Network Health in 2012. In each panel, the left figure shows the distribution (kernel densities) for switchers, while the right figure shows a bin scatter of means by decile of the HCC risk score (with confidence intervals shown in bars) to illustrate the risk-conditional distribution.

Figure A.17: Switcher vs. Stayer Cost Decomposition: Distributions and Risk Bin Scatters (cont'd)

Panel D: Inpatient Prices (multiplicative factor)

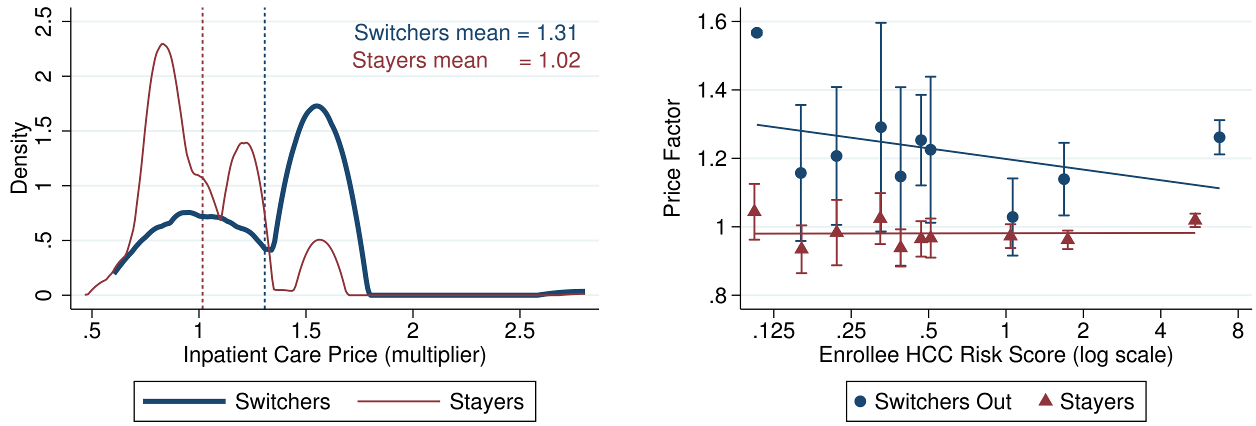

Panel E: Outpatient Prices (multiplicative factor)

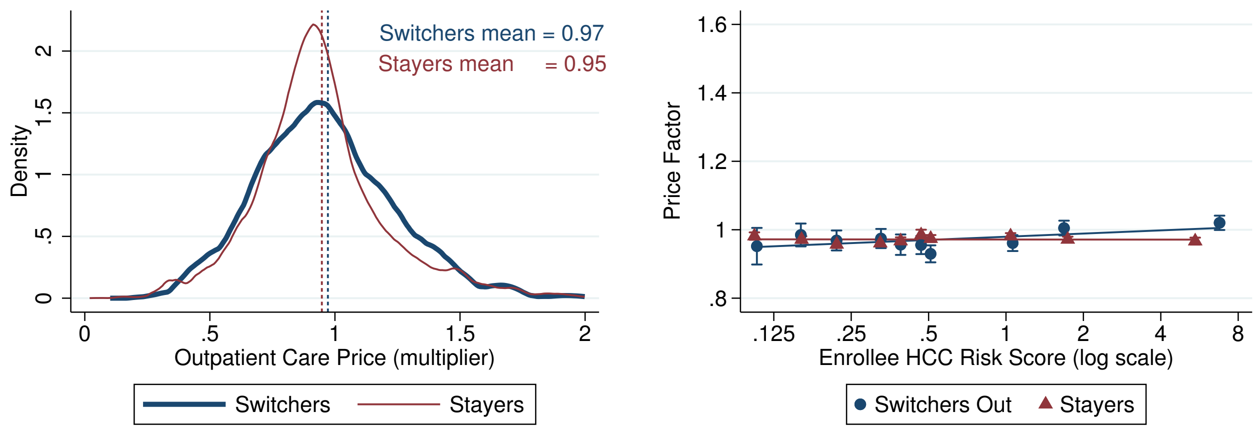

Panel F1: Share at Partners: Inpatient Care

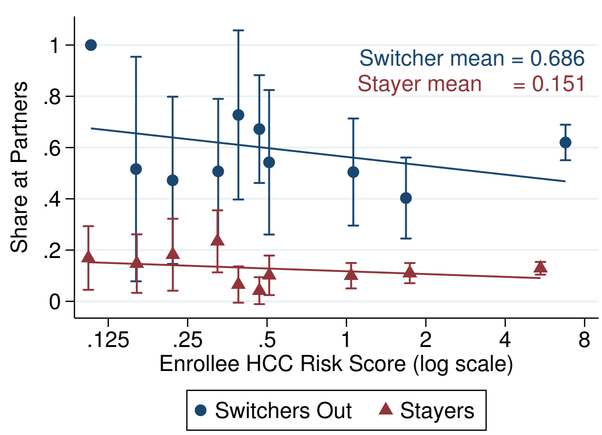

Panel F2: Share at Partners: Outpatient Care

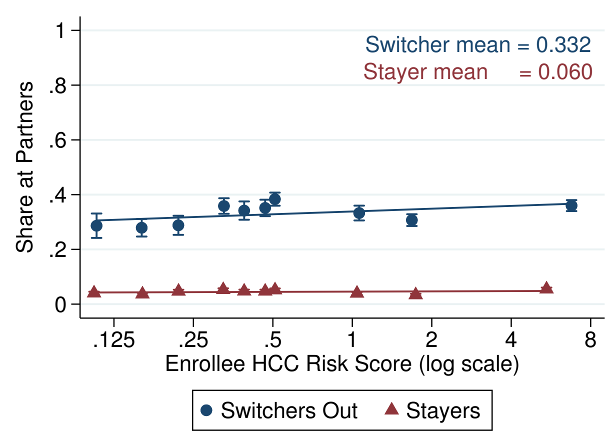

NOTE: The figures show the distribution and risk score-conditional distributions of cost components for switchers vs. stayers in Network Health in 2012. In panels D-E, the left figure shows kernel densities for switchers, while the right figure shows a bin scatter of means by decile of the HCC risk score (with confidence intervals shown in bars) to illustrate the risk-conditional distribution. Panel F shows bin scatters for inpatient (F1) and outpatient (F2) shares at Partners.

E.4 Additional Estimates of Causal Cost Effects (Moral Hazard)

Figure A.18: Event Study: Cost Reductions after 2012 Network Change

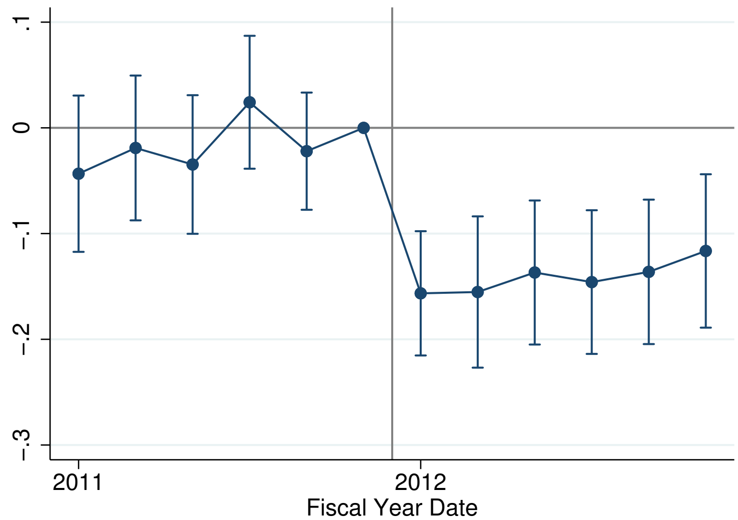

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (8) in Section 5.3 (and corresponding to Panel A of Figure 5). Estimates are from a Poisson regression with individual fixed effects and capture the cost differences 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 A.19: Event Study: Cost Reductions after 2012 Network Change, by Partners Patients

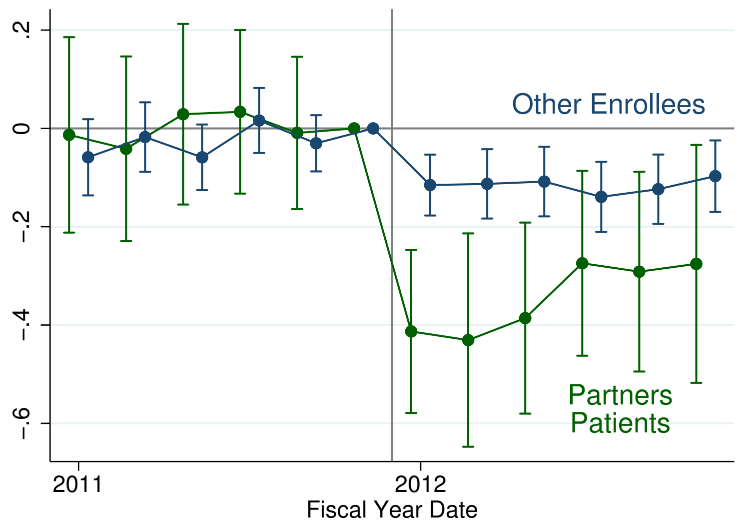

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (8) in Section 5.3, with separate interactions for γ with Partners patients (green series) versus other enrollees (blue), corresponding to Panel B of Figure 5. See note to Figure A.18 for additional information on the setup and interpretation of coefficients. This figure confirms the presence of parallel pre-trends for both groups and a share cost reduction in 2012 that is much larger for Partners patients.

Figure A.20: Cost Reductions after 2012 Network Change, by Distance to Partners Hospital

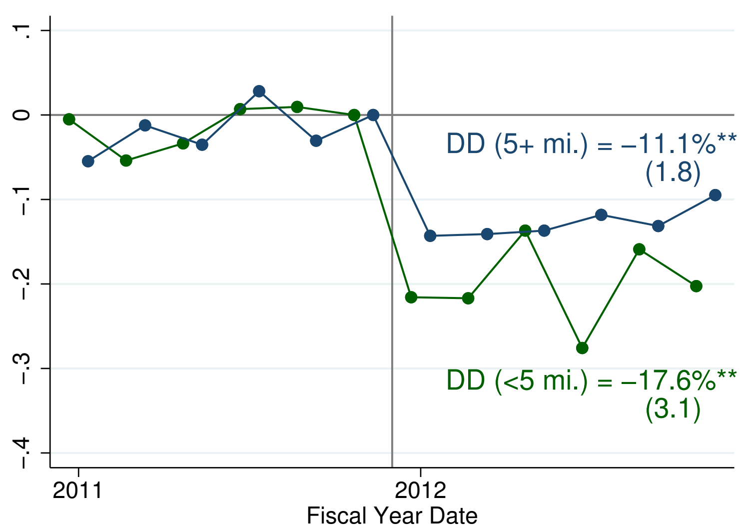

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (8) in Section 5.3, with separate interactions for γ with people living within 5 miles of a Partners hospital (green series) versus those living 5+ miles away (blue). See note to Figure A.18 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.

Figure A.21: Reductions in Quantity for Stayers after 2012 Network Change

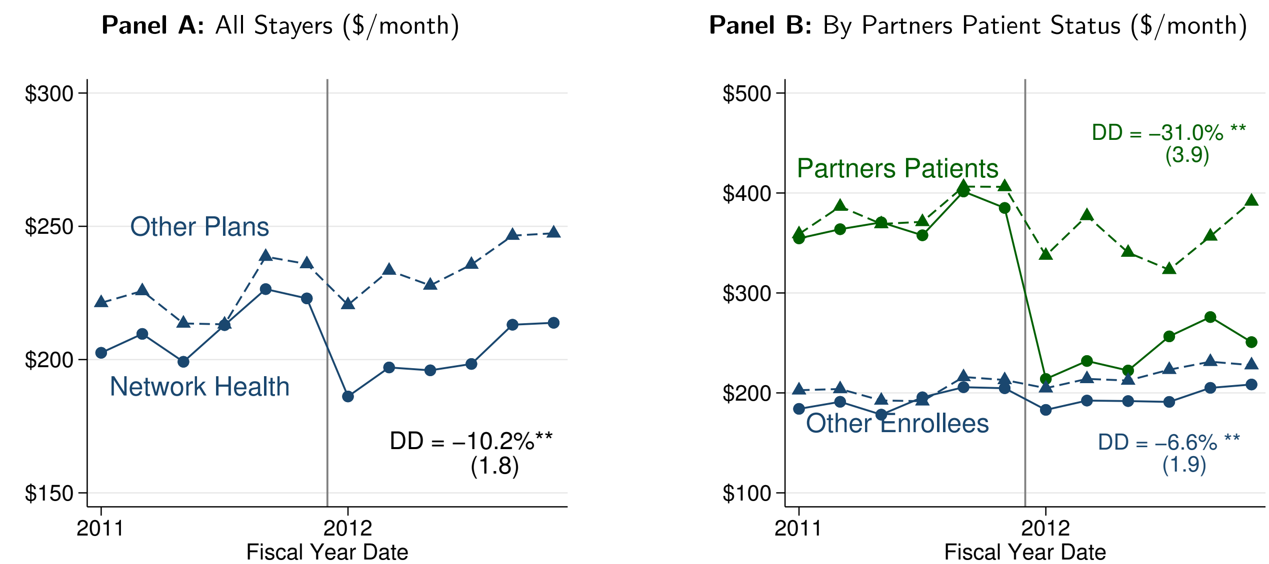

NOTE: These graphs show estimates from quantity regressions with individual fixed effects corresponding to the event study version of equation (8). The sample is "stayers" continuously enrolled in Network Health or other plans between 2011 and 2012, when Network Health narrows its network. They regression is exactly analogous to Figure 5 in the text, but with a dependent variable of quantity of care (in $ per month), rather than total costs. Quantity is defined in the price-quantity decomposition discussed in Section 5.1 and Appendix E.1. The levels of quantity are below the levels of total costs because quantity is only defined for care included in the price-quantity decomposition, which covers about two-thirds of costs.

Figure A.22: Quantity Reductions for Stayers after 2012 Network Change: Event Study

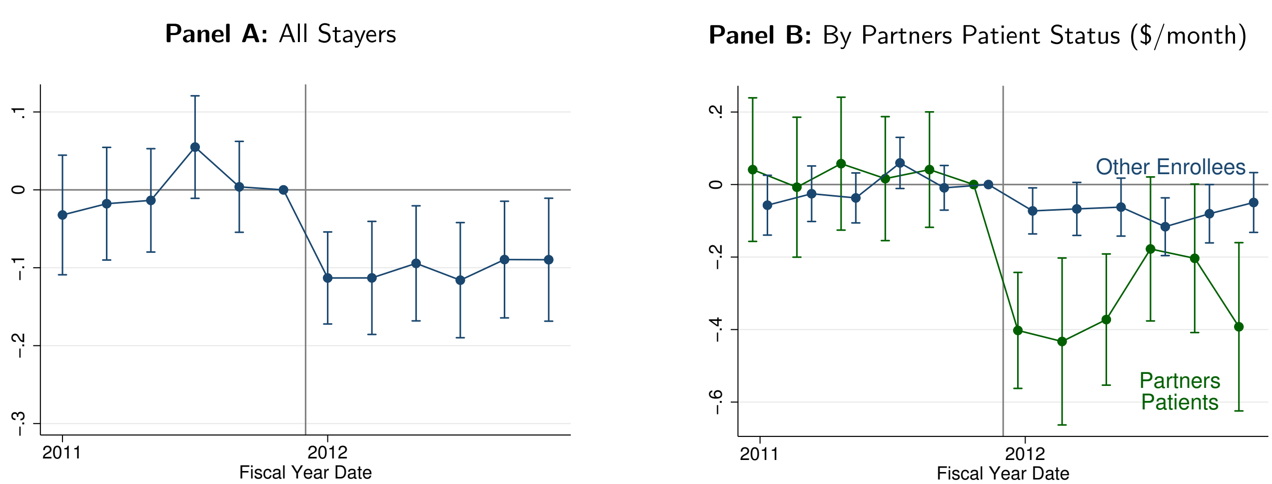

NOTE: The figure shows coefficient estimates (and 95% confidence intervals) of $\hat{\gamma}$ from the event study version of regression (8) in Section 5.3, with a dependent variable of quantity of care. Panel B has separate interactions for γ with Partners patients (green series) versus other enrollees (blue). This figure confirms the presence of parallel pre-trends for both groups and a quantity reduction in 2012 that is much larger for Partners patients.

F Appendix: Structural Model and Estimation Details

F.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:

$$U_{i,d,t,h}^{Hosp} = \underbrace{\gamma_1 \left(Z_{i,d,t}\right) \cdot Dist_{i,h}}_{\text{Distance}} + \underbrace{\gamma_2 \left(Z_{i,d,t}\right) \cdot X_h + \gamma_3 \cdot PastPatient_{i,h,t}}_{\text{Hospital characteristics x Patient observables}} + \underbrace{\eta_h}_{\text{Hospital dummy}} + \underbrace{\epsilon_{i,d,t,h}}_{\text{Logit error}} \tag{19}$$

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

$$u_{i,j,d,t,h}^{Hosp} = U_{i,d,t,h}^{Hosp} - \kappa_j \left(Z_{i,t}\right) \cdot 1\left\{h \notin N_{j,t}\right\} \tag{20}$$

The specification in (19) 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 (19) 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. 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 A.10 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.

F.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:

$$\begin{aligned} EU_{i,d,e,t,j}\left(N_{j,t}\right) &= E\left[\max\left\{\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right) + \varepsilon_{i,d,e,t,h}\right\}\right] \\ &= \log\left(\sum_h \exp\left(\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right)\right)\right) \end{aligned} \tag{21}$$

where $\hat{u}_{i,d,e,t,j,h}\left(N_{j,t}\right)$ is the utility function from (20) excluding the logit error term. (Note that I explicitly include emergency status e in the subscripts here; in equation (19) 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.[74] 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 hospitalization data for the $\{d,e\}$ and age-sex group cell.

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

Table A.10: 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).

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

$$NetworkUtil_{i,j,t}\left(N_{j,t}\right) = \sum_{d,e} freq_{i,d,e,t} \cdot EU_{i,d,e,t,j}\left(N_{j,t}\right) \tag{22}$$

The network utility in (22) 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.

F.3 Plan Choice Model Details

The plan choice model is described in Section 6.1. This appendix describes additional model details. Table A.11 below shows a summary of estimates, and Table A.12 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 model is a standard multinomial logit choice model that allows for preference heterogeneity across consumers based on observables. The choice utility specification, as reported in equation (9) is:

$$U_{i,j,t}^{Plan} = \underbrace{\alpha\left(Z_{it}\right) \cdot Prem_{i,j,t}}_{\text{Subsidized Premium}} + \underbrace{V\left(N_{j,t}; Z_{it}, \beta\right)}_{\text{Network Value}} + \underbrace{\delta\left(Z_{it}\right) \cdot 1\{CurrPlan_{i,j,t}\}}_{\text{Inertia (current enrollees)}} + \underbrace{\xi_{j,t}\left(Z_{it}\right)}_{\text{Plan dummies}} + \epsilon_{i,j,t}^{Plan}$$

where $Prem_{i,j,t}$ is the enrollee's subsidized premium, $V\left(N_{j,t}; Z_{it}, \beta\right)$ is consumer value of the provider network, $1\{CurrPlan_{i,j,t}\}$ is a dummy for current enrollees' current plan (capturing inertia), $\xi_{j,t}\left(Z_{it}\right)$ are plan dummy variables capturing unobserved quality, and $\epsilon_{i,j,t}^{Plan}$ is the type 1 extreme value error that gives shares their logit form. Coefficients on these plan characteristics are allowed to vary with consumer observables, Z_{it}. The text of Section 6.1 discusses each of these variables. Here are some additional details about each and the consumer observables their coefficients can vary with:

1. Subsidized Premiums These are observed and included directly. Premium coefficients, $\alpha\left(Z_{it}\right)$, are allowed to vary with: (1) income groups (100-150%, 150-200%, 200-250%, and 250-300% of poverty), (2) quantile of the HCC risk score (quintiles, plus an extra group for the highest 5% risk enrollees), (3) dummies for having any chronic illness and for cancer, (4) age-sex groups, and (5) immigrant status. The full list of interactions and estimates is shown in Table A.12.

Notice that unlike a standard market, premiums vary not just across plans and years (j,t) but also across consumers for a given plan-year. As discussed in the body text, insurers (who each operate a single plan) are limited to setting pre-subsidy premiums at either the plan-year-region level (from 2007-2010) or at the plan-year level (from 2011-2013). Thus, pre-subsidy premiums vary only at the plan-region-year level. The exchange applies a subsidy schedule that varies across income groups and that also affects prices differences across plans. Subsidies are set so that the lowest-price plan always costs a targeted "affordable" amount by income – e.g., in 2009-2012 this amount is \$0 per month for enrollees with incomes below 150% of poverty, \$39 for 150-200% of poverty, \$77 for 200-250% of poverty, and \$116 for 250-300% of poverty. Subsidies for higher-price plans follow a schedule that also varies across income groups and leads to variation in *premium differences* for the same plans across incomes. For enrollees in the 0-100% of poverty group, *all plans* are subsidized to be $0 – i.e., there are no premium differences. For enrollees in the 100-300% of poverty groups, higher-price plans cost more than the cheapest plan, but the gap between plans is adjusted in a "progressive" way so that premium gaps are smaller for lower-income groups and larger for higher-income groups. Appendix B.1 includes some examples of how this variation plays out.

2. Network Valuation Networks are observed and modeled using two sets of variables. The first is the "network utility" measure from the hospital choice model, described in Appendix F.2 above. The second are variables for whether the plan covers the hospitals with which the consumer has past outpatient relationships (or the share covered if there are multiple). These variables are all observed and vary across consumers and years, so identification comes from the relationship between this panel variation and consumer plan choices.

Coefficients on network utility are allowed to vary by: (1) income groups, (2) HCC risk score quantiles, and (3) dummies for having any chronic illness and for cancer. I do not vary coefficients with age-sex groups because the illness probabilities used to define network utility already vary by age-sex groups. Coefficients on coverage of hospitals with which a consumer has relationships are allowed to vary with these same three sets of characteristics, and I also further interact these coefficients with whether the hospital is a Partners hospital to allow for special loyalty to the star hospitals.

3. Inertia (current enrollees) To capture inertia, which is well known to affect health insurance choices, I include a dummy for current enrollees' current plan. Coefficients, $\delta(Z_{it})$, are allowed to vary with the the same observables as premium coefficients: (1) income groups, (2) HCC risk score quantiles, (3) chronic illness and cancer dummies, (4) age-sex groups, and (5) immigrant status. Including a lagged plan dummy allows for capturing inertia in a simple way, but the estimates may pick up both true inertia and persistent unobserved preference heterogeneity. Column (1) of Table A.11 shows a robustness check that includes only new/re-enrollees (for whom inertia is not relevant) and finds that remaining coefficient estimates are quite similar as in the full specification with current enrollees (column 2).

4. Plan dummy variables (unobserved quality) I include a large number of plan dummy variables and interactions to capture unobserved plan quality (e.g., insurer reputation) and to ensure proper identification of the premium coefficient. For each plan, I include separate dummies at the region-income group and region-year level, as well as interactions with age-sex groups and risk score deciles to allow unobserved quality to vary with medical risk. The CommCare program includes five regions (Boston, Central MA, Northern MA, Southern MA, and Western MA) and five income groups at which prices vary (0-100%, 150-200%, 200-250%, and 250-300% of poverty). After omitting empty cells where a plan is not available, there are 251 plan dummy variables/interactions in total. These are not all reported in Table A.12 due to space constraints but will be available in data output in the replication packet.

Discussion of Identification The specification of plan dummies is intended to aid in identifying the premium coefficients using only *within-plan* variation across income groups due to subsidies. Specifically, the plan-region-year dummies soak up any quality variation correlated with insurer pricing, which occurs at the plan-region-year level (or plan-year level from 2011-forward). The plan-region-income group dummies soak up any persistent plan preference differences across income groups within a region. The only remaining variation in premiums not soaked up by these dummies are *within-plan differences in premium changes* across income groups. Appendix B.2 and Figure A.2 show examples of this; see that section for further discussion.

Figure A.23: Premium Coefficient Identification: Market Shares around Price Changes

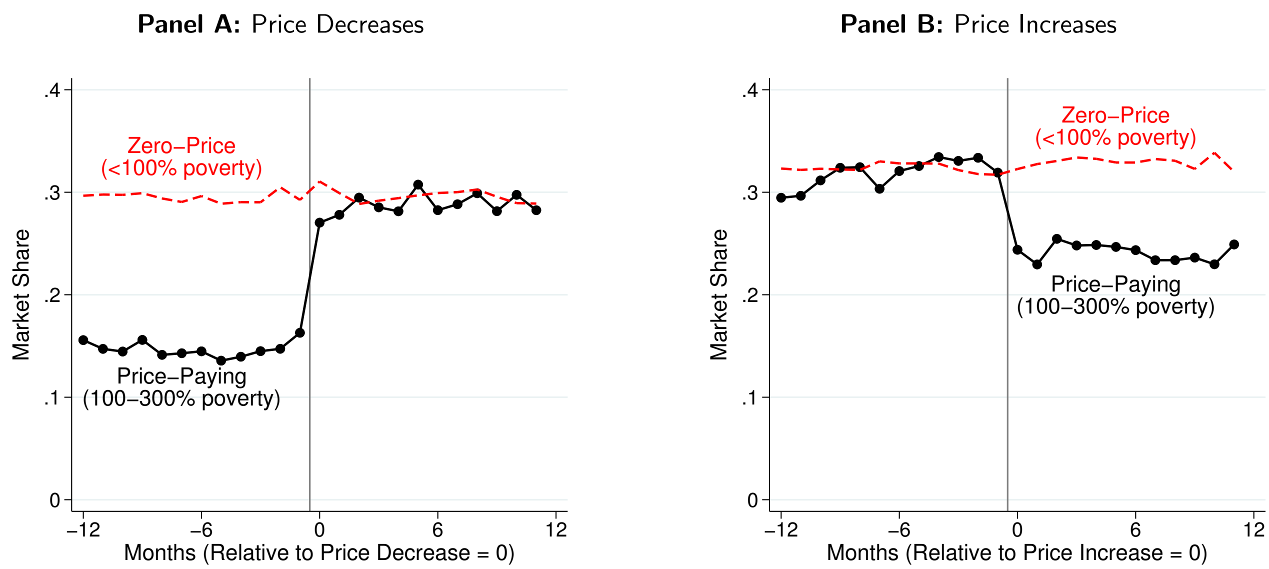

NOTE: These graphs show the source of identification for the premium coefficients in plan demand and test the key parallel trends assumption for the difference-in-differences approach. Each graph shows average monthly plan market shares among new enrollees for plans that at time 0 decreased their prices (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.

Table A.11: 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 A.12: 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)

F.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 A.13 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:

$$E\left(C_{i,j,t}\right) = \exp\left(\alpha_i + \beta_t\left(Z_i\right) + \gamma\left(Z_i\right)\cdot 1_{\{j=NH,t\geq 2012\}}\right) \tag{23}$$

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), $\beta_t\left(.\right)$ 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 (23) 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 A.13 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.

Role of Price vs. Quantity Changes 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, $dC_i = dQ_i \cdot dP_i$.[75] I can estimate regression (8) 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 $d\hat{P}_i = d\hat{C}_i/d\hat{Q}_i = \exp\left(\hat{\gamma}_C\left(Z_i\right) - \hat{\gamma}_Q\left(Z_i\right)\right)$.

Appendix Figures A.21 and A.22 show the DD estimates and event study coefficients with quantity as the outcome variable, analogous to the cost results in Figure 5 of the main text. As with costs, pre-trends are parallel, and there is a sharp quantity reduction at the start of 2012. The quantity reductions are larger in both levels and percentages for Partners patient stayers than other stayers.

Table A.13, columns (4)-(6) report estimates of the price-quantity 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

[75] To see this use the notation of Section 2 to write $C_i\left(n\right) = Q_i\left(n\right)P_i\left(n\right)$ under network n (where 0 = narrower, 1 = broader). For a narrowing of the network, $dC_i = C_i\left(0\right)/C_i\left(1\right) = \left(Q_i\left(0\right)/Q_i\left(1\right)\right)\left(P_i\left(0\right)/P_i\left(1\right)\right) \equiv dQ_i dP_i$. 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.

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.

Table A.13: Cost Model Estimates: Change in Cost with Narrower Network

Network Health x Post	**Effect on Insurer Cost**			**Decomposition**		
				Costs	Quantity	*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 (23). 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.

F.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 6B) that ΔWTP is below $\Delta Cost$. See section 6.4 for the definition of these variables. Figure A.24 replicates Figure 6B 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): This defines $\Delta Cost$ based only on changes in *quantity* of care (price-standardized utilization) due the broader network, not the effect of higher prices. This will tend to produce smaller estimates of $\Delta Cost$. Quantity reductions are estimated using the method in Appendix F.4 and specifically the estimates in column (5) of Table A.13. Figure A.24A 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 $\Delta Quantity$.[76] 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 the incremental costs of the broader network using Partners prices that are counterfactually lower, which generates a lower estimate of $\Delta Cost$. 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.[77] 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 $C_{ij}^{ALT}(1) = (1-\phi)\, C_{ij}^{Partners}(1) + C_{ij}^{Other}(1)$, where ϕ is a Partners price reduction factor. It then defines $\Delta Cost_{ij}^{ALT} = C_{ij}^{ALT}(1) - C_{ij}(0)$.[78] I consider price reductions (ϕ) of 10%, 25%, and 50%, reflecting a range of possible price reductions and/or markups.[79] Even with a 50% price reduction

[76] The decomposition excludes items like prescription drugs, inpatient rehab, and some inpatient/outpatient costs that are paid in non-standard ways (see Appendix E). 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 A.13, columns (3) vs. (4) – which is consistent with the two moving in the same direction.

[77] 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.

[78] 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.

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

(an extreme 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): This 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 (10) 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 $\Delta WTP^{ALT} = \Delta V_i / (-\widetilde{\alpha})$, where $-\widetilde{\alpha}$ is a uniform social marginal utility of money (e.g., reflecting a cost of redistribution).[80] I consider two possible values for $\widetilde{\alpha}$: (1) the average $\alpha(Z_i)$ among CommCare consumers and (2) 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 (from panel C), 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.

4. Counting only lower inpatient prices from steering patients in $\Delta Cost$ (Panel D): This panel recalculates $\Delta Cost$ by assuming that the entire cost impact of the narrow network operates through *lower inpatient hospital prices* (due to exclusion of high-price hospitals from network and associated steering to lower-price hospitals). All other cost variables – inpatient quantity, outpatient quantity and prices, and all other spending – are assumed unchanged. I estimate changes in inpatient prices by taking observed 2011 Network Health hospital admissions, re-predicting choices using the hospital choice model with the 2012 network exclusions applied, and applying the plan's hospital price estimates for 2012.[81] This modification makes a much more substantial difference, so that $\Delta Cost^{IP}$ is now less than ΔWTP across the whole distribution (it is about half as large as ΔWTP). A major reason is that inpatient costs are only about 20-25% of overall spending. Therefore, although inpatient costs fall by about 15% among the highest-WTP types (and 5-10% among lower-WTP types), the fall in *total* spending is only 1-5%. Although this is much smaller than the main estimates of $\Delta Cost$ (which includes changes in quantity and outpatient costs), a 1-4% spending reduction is consistent

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.

[80] 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.

[81] I use 2012 hospital prices because Network Health's out-of-network prices paid to Partners are much lower than its in-network 2011 prices. Its prices for all other hospitals do not change much from 2011-12.

with the estimates in Table A.13 that overall *prices* of care fall by about 3%.[82] Thus, this analysis suggests that most consumers *would be* willing to pay for star hospital coverage if the only incremental costs were via higher inpatient prices.

Figure A.24: Robustness Analysis: ΔWTP and $\Delta Cost$ of Broader Network

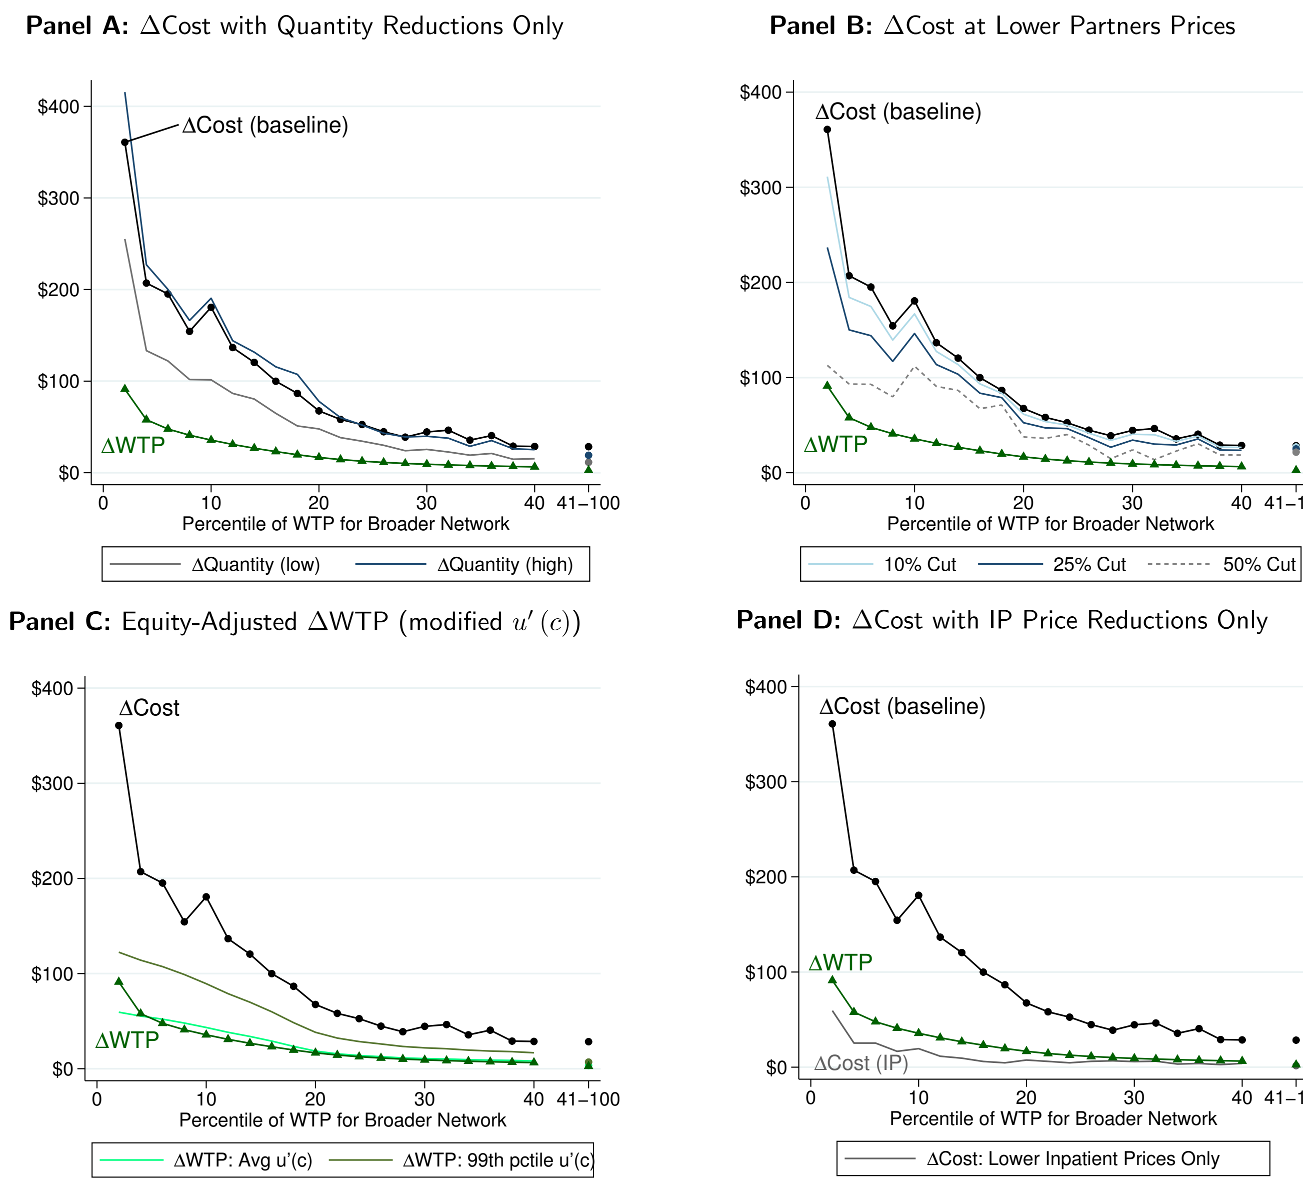

NOTE: These figures replicate Figure 6 in the body text with various modifications to $\Delta Cost$ and ΔWTP. See the note to Figure 6 and the text of Appendix F.5 for additional information describing the definition of these curves.

[82] Thus, this is consistent with the entire price reduction occurring through inpatient care, with no fall in outpatient prices. This makes sense given the finding that Partner hospitals' inpatient prices are quite high but their outpatient prices are similar to the state average.