

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

By MARK SHEPARD*

Health insurers increasingly compete on their networks of medical providers. Using data from Massachusetts’s 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 D82, G22, H75, I11, I13, I18)

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

* Kennedy School of Government, Harvard University, and NBER (email: Mark_Shepard@hks.harvard.edu). Liran Einav was the coeditor for this article. I thank 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 PhD funding support from National Institute on Aging grant T32-AG000186 (via the NBER), the Rumsfeld Foundation, and the National Science Foundation Graduate Research Fellowship. This project was approved by the IRBs of Harvard University and the NBER.

[†]Go to <https://doi.org/10.1257/aer.20201453> to visit the article page for additional materials and author disclosure statement.

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 US health care, star hospitals tend to share two features. First, they are known for advanced medical treatment and research—e.g., reflected in *US News and 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 (Coe, Lamb, and Rivera 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’s 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 nonstandard 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 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

¹ See Einav, Finkelstein, and Levin (2010) 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, Layton, and Prinz 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 versus lower cost-sharing and between health maintenance organizations (HMOs) and traditional (fee-for-service, or 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.

two nationally top-ranked hospitals (Massachusetts General Hospital [MGH] and Brigham and Women's Hospital [BWH]).

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 percent of Partners patients switch out of the plan—a more than sixfold increase relative to prior years, and strikingly high given well-known consumer inertia (Handel 2013, Ericson 2014). Relative to stayers, switchers had 108 percent higher costs, and 60 percent 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 (nonmedical) 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 I. 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.

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.

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 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 trade-off 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 *US 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 percent 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 percent of the explained variation in demand measures, versus 2–8 percent explained by “observed risk” (variables used in risk adjustment) and another 28–35 percent 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 trade-off.

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 nonstandard expensive providers dimensions. High utilization linked to medical risk explains just over half

(53 percent) 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 percent) 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 percent cost reduction for stayers, occurring through both lower prices and quantity of care. Cost reductions are much larger for Partners patients (about 30 percent, or \$175 per month) than for other enrollees (9 percent, 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 trade-offs 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, Levin, and Mahoney 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 Efficiency, however, may not be the only relevant criteria, and regulators may wish to ensure access to star hospitals based on equity or other considerations.

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 trade-offs between costs, quality, and access to top providers.

The paper proceeds as follows. Section I presents a model formalizing the paper's main ideas. Section II introduces the setting and data. Sections III to IV show reduced form evidence and analyze the mechanisms for costs. Section V presents and analyzes a structural model, and Section VI concludes.

⁴Given 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.

⁵Although 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 percent.

I. 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: (i) medical risk and (ii) costs due to use of expensive providers for care.

A. 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 versus exclude a top 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 $\varphi_i \equiv E(C_{ij}|Z_i) / \bar{C}$ 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

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

⁶A 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.

⁷This 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.

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

$$(2) \quad \Delta \pi_j = \underbrace{\sum_i [\Delta P_j - \Delta C_{ij}^{RA}] \cdot \varphi_i D_{ij}(0)}_{\text{(i) Fixed enrollment: premium and cost change}} + \underbrace{\sum_i [P_j(1) - C_{ij}^{RA}(1)] \cdot \varphi_i \Delta D_{ij}}_{\text{(ii) Selection: profitability of marginal enrollees}}.$$

Term (i) represents the impact of the plan's premium change and incremental costs (moral hazard) due to the broader network, holding enrollment fixed. Term (ii) 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 $P_j(1)$.

B. 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 $d \in \{1, \dots, D\}$, 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

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

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., $R_i/\varphi_i \rightarrow 1$). 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

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

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: (i) unobserved medical risk, R_i/φ_i , and (ii) 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 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 nonmedical 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 trade-offs, not a mere technical fix. Fundamentally, selection on moral hazard complicates risk adjustment because it becomes tied up with fundamental economic cost-quality trade-offs (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)

⁸My empirical work (Sections IV and V) 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 versus provider prices.

⁹The 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.

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 it 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.

II. Massachusetts Exchange Setting and Data

A. Setting: Massachusetts Subsidized Exchange (*CommCare*)

I study Massachusetts’s subsidized health insurance exchange—called Commonwealth Care, or CommCare. Created in the state’s 2006 “Romneycare” health reform, CommCare operated from 2006 to 2013 to provide subsidized coverage to low-income people (below 300 percent of poverty) not eligible for employer insurance or other public programs.¹⁰ Enrollees could choose among competing private plans in a centralized marketplace. Over the 2010–2013 period I focus on, the exchange featured five competing insurers and averaged 170,000 enrollees—making it a substantial market but still only a small portion of the state’s population of 6.6 million.

CommCare is a good setting to study the selection implications of provider networks (and star hospital coverage in particular) for several reasons. First, the exchange standardized essentially all benefits other than networks. By rule, each insurer offered a single plan with state-specified covered services and patient cost sharing rules.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 Health and Human Services-Hierarchical Condition Categories [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.

¹⁰ A separate market called “CommChoice” offered unsubsidized plans for all others (for research on CommChoice, see Ericson and Stane 2015a, 2015b, 2016). In the ACA, unsubsidized and subsidized enrollees are pooled into a single exchange, while people below 138 percent of poverty are eligible for Medicaid in states that have chosen to expand the program.

¹¹ The only exceptions to this identical coverage were (i) prescription drug formularies for above-poverty enrollees, subject to minimum standards, and (ii) a few “extra benefits” like gym memberships.

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

TABLE 1—HOSPITAL PRICES: MOST EXPENSIVE HOSPITALS FOR COMM-CARE INSURERS

Hospital	System	Teaching status	Raw data	Hospital price model		
			Average insurer payment (1)	Relative price		Relative patient severity (4)
				Estimate (2)	SE (3)	
1. Brigham and Women's Hospital	Partners	AMC	\$23,525	1.62	(0.04)	1.37
2. Massachusetts General Hospital	Partners	AMC	\$21,090	1.58	(0.04)	1.25
3. Boston Medical Center	BMC	AMC	\$16,478	1.29	(0.03)	1.20
4. Baystate Medical Center	Baystate	Teaching	\$13,411	1.27	(0.03)	0.99
5. UMass Medical Center	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 Medical Center	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
Nontop ten hospitals	—	—	\$7,972	0.84	—	0.88

Notes: The table shows the ten 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 to 3 show the in-network relative price estimates (for $t = 2011$) and standard errors from inpatient price model (see Section IVA), 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.)

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: MGH and BWH, which are owned by the Partners Healthcare System. *US News and World Report* perennially ranks these as the top two hospitals statewide and among the top ten 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, Massachusetts CHIA 2014a). 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 IVA. The table reports inpatient price estimates for the ten 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 (versus 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 percent above the next-highest hospital.¹³ The star

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 (Massachusetts CHIA 2014b), BWH, and MGH have the highest casemix-adjusted costs of any large general acute hospital, with costs about 30–50 percent 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 percent) than costs (30–50 percent), suggesting that markups are also high at the star hospitals.

hospitals also treat some of the sickest patients, with average severities 25–37 percent 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 ten 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 percent 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 (presubsidy) 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 online Appendix B.

B. Administrative Data: Plan Enrollment and Insurer Claims

I use administrative data on enrollment and insurance claims for all CommCare plans and enrollees from fiscal years 2007 to 2014 (Massachusetts Health Connector 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 has multiple claims per admission—is an involved process; I discuss details in online 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 Annual Survey (AHA 2014) and define travel distance using the driving distance from the patient’s zip code centroid to each hospital (Google Maps 2014).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

¹⁴The data were obtained via a data use agreement with the Massachusetts Health Connector (2014), the exchange regulator. To protect enrollees’ privacy, the data were purged of all identifying variables. This data is supplemented with public records on plan hospital networks (Massachusetts Health Connector 2013).

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

is constructed at the level of instances of enrollees making a plan choice, which occur at two times: (i) when an individual newly enrolls in CommCare (or re-enrolls after a gap), and (ii) 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 $\times$ 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 versus quantities, as discussed in Section IVA.

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 online 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 using the name of the billing provider on the claims. The key variables for my analysis are whether enrollees received nonemergency department 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.—Online 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 percent 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 percent of the market) in steady state—giving me a significant population of active choosers for plan demand estimation.

C. 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 online 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, the exchange introduced new rules encouraging insurers to compete more aggressively

¹⁶This 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.

¹⁷These 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 full Boston network.




FIGURE 1. CHANGES FOR NETWORK HEALTH AROUND 2012 NETWORK CHANGE

Notes: The figures show average enrollee cost per month (panel A) and Partners hospital use shares (panel B) by enrollees in Network Health and all other CommCare plans. Each point is a bimonthly or 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 percent), 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.

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 percent, 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 percent. 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 reflect 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 percent of consumers who left Network

¹⁸There 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 percent 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.

¹⁹These 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.

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, CeliCare 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

III. 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 I. To do so, I study the natural experiment created by Network Health dropping the star Partners hospitals in 2012, as just described in Section IIC. 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 IIIA shows the main evidence from plan switching choices in 2012, and Section IIIB examines the role of sickness and preferences in explaining switching choices.

A. 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: (i) through plan switching by current enrollees and (ii) 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

²⁰In testimony to the Massachusetts Health Policy Commission (HPC 2014), CeliCare's CEO wrote: "For the contract year 2012, Network Health Plan removed Partners hospital system and their PCPs [primary care physicians] from their covered network. As a result, the CeliCare membership with a Partners PCP increased 57.9 percent. CeliCare'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 and MGH PCPs as of July 1, 2013."

²¹Moreover, 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.

²²The 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.




FIGURE 2. SPIKE IN SWITCHING RATES OUT OF NETWORK HEALTH AT 2012 NETWORK NARROWING

Notes: 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 nonemergency room outpatient care).

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 online 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 percent, more than four times the 2.4 percent rate in 2010–2011. 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 percent for people within 5 miles of Partners, versus a steady 5 percent rate for those more than 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 percent—a more than *twentyfold* increase over the rate for the prior two years (2.1 percent). Switching also jumps to 24 percent for patients of other dropped hospitals (versus 1.7 percent in the prior two years). By contrast, switching for all other enrollees was much lower (3 percent) 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 percent higher than stayers (\$675 versus \$324 per month). CommCare's risk adjustment narrows this cost gap to 60 percent (\$508 versus \$317 per month) but does not close it. Indeed, the

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 percent of Network Health enrollees in 2011 but (because they are so much more likely to leave) comprise 67 percent of switchers out. Other dropped hospitals' patients represent 8 percent of 2011 enrollees but 17 percent of switchers out. Thus, these two groups together comprise the vast majority (84 percent) of switchers out in 2012.




FIGURE 3. ADVERSE SELECTION EVIDENCE: AVERAGE COST OF SWITCHERS AND STAYERS (\$ PER MONTH)

Notes: 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.

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 percent (20 percent) 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 online 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.

²⁴Consistent with this interpretation, 91 percent of the 2012 switchers (and 98 percent of Partners patient switchers) shift to one of the two plans that still covers Partners (CeltiCare and NHP).

Robustness Checks.—In online Appendix C.1, I implement three analyses to check the robustness of these adverse selection findings: (i) studying switching by zero-premium enrollees, for whom there is no concurrent change in Network Health’s premium along with the narrower network; (ii) examining new enrollee choices, which are not subject to inertia; and (iii) 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, and 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.

B. Role of Sickness and Preferences

The theory in Section I emphasizes two channels for costs and demand for the star hospitals: sickness and (nonmedical) 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 to 2012,²⁵ I estimate logit regressions of the form

$$(5) \quad \text{SwitchPlans}_{i,t} = \text{logit}(\alpha + \beta \cdot X_{i,t} + \gamma \cdot X_{i,t} \mathbf{1}_{\{t=2012\}}),$$

where $X_{i,t}$ are enrollee characteristics (e.g., distance to Partners). In the regression, α and β capture plan switching patterns in 2009–2011, and γ captures the *excess* switching in 2012.

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 more than 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 percent of enrollees (odds ratio = 9.72). The third panel shows that a similar relationship holds for *unobserved* sickness, defined as the ratio of the HCC risk score to the risk score used by CommCare.

²⁵I exclude 2013–2014 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 2013–2014.

²⁶For 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.

²⁷Results are similar with a multivariate regression; see online Appendix Figure A.13. Online 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 4. SPIKE IN PLAN SWITCHING IN 2012: ROLE OF SICKNESS AND PREFERENCES (DISTANCE)

Notes: 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 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 percent and sickest 5 percent of enrollees based on HCC risk score.

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 percent 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 more than 25 miles away, the sickest 5 percent 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).

These results suggest that *both* sickness and nonmedical 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: (i) switching plans in 2012, and (ii) being a Partners patient in 2011. Online Appendix D.1 discusses the method details.²⁸

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.

Overall, the estimates suggest 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 percent 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 percent of variation, and a much richer set of “unobserved risk” measures derived from the claims explains another 28–35 percent. 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 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). Online Appendix D.2 discusses these two channels further and provides some evidence that both are involved.

IV. Understanding Costs Driving Adverse Selection

The evidence in Section III 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 versus quantity in Sections IVA and IVB. 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 IVC shows evidence of this prediction.

A. Decomposition of the Two Cost Dimensions

The theory in Section I shows that consumers may incur high costs through two dimensions: (i) medical risk (the standard channel) and (ii) use of expensive providers (the nonstandard 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 ($\kappa_{ij}(\cdot)$). 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}) versus 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 IVC below.

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

- (i) Costs due to *high provider prices* reflect the expensive provider channel.
- (ii) 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.
- (iii) *Causal changes in quantity and costs due to the network change* reflect the expensive provider use channel.

Sections IVB and IVC 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 versus 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 $s \in \{1, \dots, S\}$ (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

²⁹The 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 FFS. See online Appendix E for further details.

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

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

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 versus 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 (Healthcare Common Procedure Coding System), 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 versus nonfacility) and type of care (e.g., medical versus behavioral health versus dental care). For inpatient care, the service unit is an admission for a particular 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). Online 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 $\hat{Q}_{it}^{risk} = E[Q_{it}|Z_{it}] = f(Z_{it}; \hat{\theta})$, where $f(\cdot; \theta)$ is the two-part model's prediction function (see online Appendix E.1). I implement this using two sets of Z_{it} variables: (i) only “observed risk” variables included in risk adjustment (age and CommCare's risk score), and (ii) a broader set of variables from the claims (including diagnoses and the concurrent HCC risk score). After estimating $\hat{Q}_{it}^{risk}$, “residual quantity” is defined as the remaining factor explaining quantity: $\hat{Q}_{it}^{resid} \equiv Q_{it}/\hat{Q}_{it}^{risk}$.

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:³¹

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

This equation lets me decompose the share of group cost differences (e.g., stayers versus switchers in 2012) that are driven by (i) risk-predictable quantity, (ii) residual

³⁰This 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 percent of enrollee-years for outpatient and overall costs, though only 4 percent 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.

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

quantity, and (iii) provider prices. Its multiplicative form suggests decomposing log differences for each factor, which are additive.

B. Cost Decomposition Results

I now apply the method just outlined to decompose cost differences between switchers out versus 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.

Online Appendix Figures A.16 and A.17 show descriptive plots of components of the decomposition for switchers versus stayers, both overall and conditional on enrollee risk score. The patterns suggest that switchers are high cost on nearly all metrics. Switchers have (i) higher risk-predictable quantity using either “observed risk” factors (used by CommCare) or all risk measures, (ii) higher residual quantity conditional on risk, and (iii) 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 percent) for inpatient and 1.96 (or +96 percent) for outpatient costs. For the two combined, stayers have 126 percent higher costs—similar to the 108 percent 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 percent of cost differences, while adding a broader set of risk variables explains 53–55 percent 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 versus 4 indicates that “unobserved risk”—not captured by CommCare’s risk adjustment but predictable using concurrent risk measures – explains 29–37 percent of adverse selection.

The second key finding is that provider prices (column 6) explain a meaningful 22 percent share of inpatient cost differences, though only 4 percent 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’

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

TABLE 2—DECOMPOSITION OF SWITCHERS' HIGH COSTS: PRICE VERSUS QUANTITY

	Spending (\$/month)	Quantity of care			Residual factor	Provider price factor
		Overall quantity	Predicted by risk vars.			
			Used for risk adj.	All risk variables		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Inpatient care</i>						
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)
Percent of log diff. explained (percent)	100	78	18	55	24	22
<i>Panel B. Outpatient care</i>						
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)
Percent of log diff. explained (percent)	100	96	24	53	43	4
<i>Panel C. Combined (IP + OP care)</i>						
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)
Percent of log diff. explained (percent)	100	88	21	53	35	12

Notes: 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 IVA. For inpatient costs (IP; panel A), outpatient costs (OP, panel B), and the sum of the two (panel C), the columns decompose switcher-stayer differences into components: overall quantity (column 2), risk-predictable quantity (columns 3 and 4); residual quantity (column 5); and provider prices (column 6). Columns 3 and 4 differ in the risk covariates used. Column 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). Column 4 adds concurrent variables observed in 2011: diagnoses and a spline in the HCC risk score.

over four times higher propensity to choose Partners hospitals (69 percent share for switchers versus 15 percent for stayers), whose inpatient prices are 45 percent above average. For outpatient care, switchers are also much more likely to choose Partners (33 percent versus 6 percent share), but interestingly Partners’ outpatient prices are not high (they are within 3 percent of the statewide average).

The third finding in Table 2 is that residual quantity (column 5) explains a substantial share of cost differences: 24 percent for inpatient and 43 percent for outpatient care. As noted, this component is more challenging to interpret. It may reflect either further unobserved risk *or* provider impacts on treatment intensity. To gain additional insight, online Appendix Table A.9 examines how this residual quantity relates to propensity to use Partners, both using the raw ordinary least squares relationship and using distance as an instrumental variable. Partners patients

for risk. This is consistent with the findings in Section IIIB that observable sickness explains only a small share of the demand for the star hospitals.

consistently have high residual quantity, with levels about 20 percent higher in the instrumental variable 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.

C. Evidence from Causal Cost Changes due to Narrower Network

My model in Section I 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 VB.

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 to 2012, relative to a control group of stayers in other plans. Limiting the sample to stayers and the 2011–2012 period, I run a Poisson regression with individual and time fixed effects. The specification is

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

where $C_{i,j,t}$ is insurer cost on individual i at time t , α_i is an enrollee fixed effect, $\beta_t(\cdot)$ 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 (DD) but in a nonlinear 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 $\gamma(\cdot)$ 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 versus other plans (no Z_i heterogeneity). To visualize levels along with changes, I report the predicted means for Network Health ($= \exp(\bar{\alpha}_{NH} + \beta_t + \gamma_t)$) and for other plans ($= \exp(\bar{\alpha}_{Oth} + \beta_t)$), where the $\bar{\alpha}_g$ 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 percent reduction (standard error = 1.6 percent), or about \$45 per month. By

³³I 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 versus quantity. However, all main results are robust to using a linear fixed effects specification.




FIGURE 5. COST REDUCTIONS FOR STAYERS AFTER 2012 NETWORK CHANGE

Notes: 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 $\exp(\bar{\alpha}_{Oth} + \beta_t)$ (for other plans) and $\exp(\bar{\alpha}_{NH} + \beta_t + \gamma_t)$ (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) versus all other enrollees (in blue), with solid lines continuing to denote Network Health and dashed lines other plans.

contrast, costs for other plans change very little and move in parallel to Network Health’s costs aside from the one-time fall at the start of 2012.³⁴

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 versus all other enrollees, the strongest predictor of selection. The graph shows two facts. First, Partners patients are much higher cost in the preperiod (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 versus −\$30 per month) and in percentage terms (−30.6 percent versus −9.2 percent). Online Appendix Figure A.19 plots the γ_t estimates, confirming the presence of parallel pretrends 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 percent in 2011 (\$619 versus \$285 per month) down to +40 percent in 2012 (\$406 versus \$290).³⁵

³⁴ Online 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 5 panel B). 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 percent (standard

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

Online Appendix F.4 shows how this approach can be used to further decompose the causal effects into changes in quantity versus price of care, following the decomposition in Section IVA. 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 percent 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).

V. Policy Analysis and Welfare Trade-Offs

A. 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: (i) new enrollments in the exchange (including re-enrollments after a break) and (ii) plan switching decisions at annual open enrollment. For consumer i choosing at time t , the utility for plan j equals

$$(9) \quad U_{i,j,t}^{Plan} = \underbrace{\alpha(Z_{it}) \cdot \text{Prem}_{i,j,t}}_{\text{Subsidized premium}} + \underbrace{V(N_{j,t}; Z_{it}, \beta)}_{\text{Network value}} \\ + \underbrace{\delta(Z_{it}) \cdot \mathbf{1}\{\text{CurrPlan}_{i,j,t}\}}_{\text{Inertia (current enrollees)}} + \underbrace{\xi_{j,t}(Z_{it})}_{\text{Plan dummies}} + \epsilon_{i,j,t}^{Plan}.$$

error = 3.1 percent), compared to a smaller fall for further enrollees of 11.1 percent (standard error = 1.8 percent); see online 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 online 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 percent versus the 40 percent 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, online 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.

In addition to the “logit” error ($\epsilon_{i,j,t}^{Plan}$), plan utility depends on four plan characteristics: (i) subsidized premiums, (ii) provider networks, (iii) inertia for current enrollees in their current plan, and (iv) 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 percent). Online Appendix 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: (i) because of insurer pricing, which occurs at the plan-year level, (in some years, separately across five regions); and (ii) 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 $V(\cdot)$. First, I follow past work (starting with Capps, Dranove, and Satterthwaite 2003) by including a “network utility” measure derived from an estimated hospital choice model. Online Appendix F.1 and 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.

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.

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 (presubsidy) plan prices at either the region-year level (prior to 2011) or at the yearly level (2011 onward). Insurers by rule may not vary prices at a more detailed 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 and 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 DD in a nonlinear 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. Online 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. Online Appendix Tables A.11 and 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.

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. Online 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

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 percent. 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.

There is also substantial inertia in consumers' plan switching decisions, with the average coefficient of 4.413 (standard error = 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 percent 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 (standard error = 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 (standard error = 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 (standard error = 0.021), or \$19.43 per month. As in the reduced form evidence (Figure 2 panel B), 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 VD below.

B. 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 I, the goal is to estimate $E(C_{ijt}(0)|i \in G)$ and $E(C_{ijt}(1)|i \in G)$ 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 2011–2012 as it narrows

network value are robust to any challenges in distinguishing inertia versus unobserved preferences. I therefore use column 2 for the remainder of the analysis.

⁴²This 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 percent increase relative to the average plan price of about \$400. The typical firm-perspective elasticity is therefore about -10.4 ($= -26.1$ percent share change/2.5 percent price change).

⁴³Converting 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).

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: (i) estimating expected costs under the plan's *observed* network ($n = 1$ in 2011 and $n = 0$ in 2012), and (ii) 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 $C_{ijt}^{obs}(1)$). Assume that realized costs equal expected costs ($C_{ijt}(1)$) plus an idiosyncratic shock: $C_{ijt}^{obs}(1) = C_{ijt}(1) + \epsilon_{ijt}$. If the variables defining group G are known at the time when expected costs are defined,⁴⁴ then $E(\epsilon_{ijt}|i \in G) = 0$ and expected costs for group G can be estimated as the average of realized costs: $\bar{C}_{G,t}(1) \equiv \frac{1}{N_G} \sum_{i \in G} C_{ijt}^{obs}(1) \rightarrow E(C_{ijt}(1)|i \in G)$ 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 IVC, which are identified from stayers in Network Health from 2011 to 2012, relative to a control group of stayers in other plans. The identification is based on a DD logic, and Figure 5 shows evidence of parallel pretrends. 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 $\gamma(\cdot)$ comes from the reverse experiment of a narrower network. Online Appendix Table A.13 shows the results, with columns 3–6 reporting estimates for insurer cost, quantity, and prices. Given an estimate of either $\bar{C}_{G,t}(1)$ or $\bar{C}_{G,t}(0)$ 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 IVC 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 *underestimate* of $\Delta Cost$, which would reinforce the finding that these are larger than consumer WTP for star hospital coverage.

⁴⁴This 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.

TABLE 3—ANALYSIS OF INCENTIVES FOR 2012 NETWORK NARROWING

	Observed values		Breakdown of change in profit			
	2011 broad network	2012 narrow network	Selection incentive		Fixed enrollment	
	(1)	(2)	Switch out (3)	Switch in (4)	Δ Price (5)	Δ Cost (6)
<i>Premium and medical costs (\$/month)</i>						
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		
<i>Demand and profits</i>						
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				
Δ Profit 2011–2012		+\$0.65	+\$0.96	+\$0.45	−\$2.56	+\$1.80

Notes: 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 and 2) into selection incentives (columns 3 and 4) vs. fixed enrollment price/cost changes (columns 5 and 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 and 2 show premiums, costs, demand, and profits (before admin costs) directly from the data. Columns 3 and 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 to 2012. Columns 5 and 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 "base-line" costs under the narrow network ($C^{RA}(0)$) and incremental costs of the broader network (ΔC^{RA}).

C. Role of Selection Incentive in 2012 Network Change

I can use the model to break down the role of selection versus 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 and 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 and 4) and price/cost changes with fixed enrollment (columns 5 and 6).

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 percent) in raw

⁴⁵I 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.

⁴⁶Two 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 online 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).

terms, and by \$83 after risk adjustment (21 percent). Therefore, its profit margin increases by \$21 per month (38 percent), outweighing a modest decline in demand and leading to \$0.65 million higher profits.

Columns 3 and 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 percent of baseline 2011 profits, and 78 percent 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 (column 5) exceed the cost savings (column 6).

The table also illustrates the interaction of selection and moral hazard, as suggested by Section IVC. 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 percent). By contrast switchers in have both lower baseline (\$228) and incremental costs (\$32, or 14 percent). This pattern of selection on moral hazard contributes to the challenges of risk adjustment and the difficult welfare/policy trade-offs involved, which I discuss next.

D. Analysis: WTP and Cost Curves for Star Hospital Coverage

I next analyze WTP and costs under the broader 2011 versus narrower 2012 networks in the style of Einav, Finkelstein, and Cullen (2010; henceforth, 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):

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

where $V(\cdot)$ 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

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 percent of the 3,207 hospital beds in the dropped hospitals. Partners patients comprise 67 percent of the switchers out of Network Health in 2012 (versus 8 percent patients at other dropped hospitals).

⁴⁸ I do not have $\alpha(Z_i)$ estimates for below-poverty enrollees, so for them I use the α estimates for comparable 100–150 percent 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

Notes: These graphs show cost and 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).

from Section VB. 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 (dashed light-blue curve) makes a large difference, but costs are still steeply downward sloping for the broad network. Risk-adjusted costs in the top 2 percent WTP bin are \$628 per month, about 50 percent larger than at the twentieth percentile (\$416) and twice the cost at the fortieth percentile (\$309). Second, and by contrast, risk-adjusted costs under the *narrower network* (dashed red curve) are much flatter. Except for the top 2 percent 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 6 panel B 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 three to six times throughout the distribution. Thus, under a standard surplus measure, the broader network with the star hospitals is 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 percent of WTP place substantially higher value on the network—about \$90 per month, or almost twice the average enrollee premium in CommCare⁴⁹—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’ WTP appears fairly robust. Online Appendix F.5 presents robustness checks on ΔWTP and $\Delta Cost$ estimates, including (i) counting only quantity reductions in $\Delta Cost$, (ii) recalculating $\Delta Cost$ using 10–50 percent lower Partners prices (reflecting possible markups above true marginal costs), (iii) redefining ΔWTP based on a lower social marginal utility of money, and (iv) 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 (iv)), $\Delta Cost$ 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 overuse.

⁴⁹For further context, Finkelstein, Hendren, and Shepard (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 (2015a) study a higher-income Massachusetts population and find that *typical* WTP for a broad network (that includes Partners) versus a narrower network (that excludes Partners) is between \$68 and 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.

⁵⁰A 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, the ΔWTP curve would be understated. State dependence (e.g., due to a cost of switching doctors), which I analyze further in online 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.

VI. 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's 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 options—star hospitals in my study, but the same idea could apply to any expensive provider, drug, or treatment—can naturally lead to adverse selection, and specifically selection on moral hazard. Selection on moral hazard is not just an empirical curiosity but affects welfare and policy implications. Typically, economists think of adverse selection as leading to too little access to (or enrollment in) generous insurance, creating a rationale for mandates or subsidies. But selection on moral hazard complicates the analysis because people with the greatest demand for a generous benefit also have the largest cost increases from it. This poses a challenge for standard risk adjustment (Einav et al. 2016) and can make consumer sorting inefficient with any single pooled premium (Bundorf, Levin, and Mahoney 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

- American Hospital Association (AHA).** 2014. "American Hospital Association Annual Survey, 2007-2013." Accessed via DUA with the AHA through the National Bureau of Economic Research.
- Azevedo, Eduardo M., and Daniel Gottlieb.** 2017. "Perfect Competition in Markets with Adverse Selection." *Econometrica* 85 (1): 67-105.
- Breyer, Friedrich, M. Kate Bundorf, and Mark V. Pauly.** 2011. "Health Care Spending Risk, Health Insurance, and Payment to Health Plans." In *Handbook of Health Economics*, Vol. 2, 691-762. Amsterdam: Elsevier.
- Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad.** 2017. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *Quarterly Journal of Economics* 132 (3): 1261-1318.
- Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston.** 2014. "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program." *American Economic Review* 104 (10): 3335-64.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney.** 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review* 102 (7): 3214-48.
- Capps, Cory, David Dranove, and Mark Satterthwaite.** 2003. "Competition and Market Power in Option Demand Markets." *RAND Journal of Economics* 34 (4): 737-63.
- Carey, Colleen.** 2017. "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D." *American Economic Journal: Economic Policy* 9 (1): 38-73.
- Chiappori, Pierre-André, and Bernard Salanie.** 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy* 108 (1): 56-78.
- Clemens, Jeffrey, and Joshua D. Gottlieb.** 2017. "In the Shadow of a Giant: Medicare's Influence on Private Physician Payments." *Journal of Political Economy* 125 (1): 1-39.
- Coakley, Martha.** 2010. *Massachusetts Attorney General Investigation of Health Care Cost Trends and Cost Drivers: Preliminary Report*. Boston: Office of Mass. Attorney General.
- Coe, Erica Hutchins, Jessica Lamb, and Suzanne Rivera.** 2017. *Hospital Networks: Perspective from Four Years of the Individual Market Exchanges*. Pittsburgh, PA: McKinsey and Company.
- Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen.** 2019. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." *Quarterly Journal of Economics* 134 (1): 51-107.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya.** 2021. "Can Health Insurance Competition Work? Evidence From Medicare Advantage." *Journal of Political Economy* 129 (2): 570-606.
- Decarolis, Francesco, and Andrea Guglielmo.** 2017. "Insurers Response to Selection Risk: Evidence from Medicare Enrollment Reforms." *Journal of Health Economics* 56: 383-96.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877-921.
- Einav, Liran, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf.** 2016. "Beyond Statistics: The Economic Content of Risk Scores." *American Economic Journal: Applied Economics* 8 (2): 195-224.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin.** 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics* 2 (1): 311-36.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen.** 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review* 103 (1): 178-219.
- Einav, Liran, Mark Jenkins, and Jonathan Levin.** 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica* 80 (4): 1387-1432.
- Ericson, Keith M. Marzilli.** 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy* 6 (1): 38-64.
- Ericson, Keith M. Marzilli, and Amanda Starc.** 2015a. "Measuring Consumer Valuation of Limited Provider Networks." *American Economic Review* 105 (5): 115-19.
- Ericson, Keith M. Marzilli, and Amanda Starc.** 2015b. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics* 97 (3): 667-82.

- Ericson, Keith M. Marzilli, and Amanda Starc.** 2016. "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange." *Journal of Health Economics* 50: 71–85.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard.** 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530–67.
- Geruso, Michael, and Timothy J. Layton.** 2017. "Selection in Health Insurance Markets and Its Policy Remedies." *Journal of Economic Perspectives* 31 (4): 23–50.
- Geruso, Michael, Timothy Layton, and Daniel Prinz.** 2019. "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges." *American Economic Journal: Economic Policy* 11 (2): 64–107.
- Geruso, Michael, and Thomas G. McGuire.** 2016. "Tradeoffs in the Design of Health Plan Payment Systems: Fit, Power and Balance." *Journal of Health Economics* 47: 1–19.
- Ghili, Soheil.** Forthcoming. "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance." *Marketing Science*.
- Glied, Sherry.** 2000. "Managed Care." In *Handbook of Health Economics*, Vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 707–53. Amsterdam: Elsevier.
- Google Maps.** 2014. "GoogleMaps API Driving Distance Dataset." Shared with author by Prof. Amanda Starc, who accessed it by calling Google Maps API.
- Gruber, Jonathan, and Robin McKnight.** 2016. "Controlling Health Care Costs through Limited Network Insurance Plans: Evidence from Massachusetts State Employees." *American Economic Journal: Economic Policy* 8 (2): 219–50.
- Handel, Benjamin R.** 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.
- Ho, Kate, and Robin S. Lee.** 2017. "Insurer Competition in Health Care Markets." *Econometrica* 85 (2): 379–417.
- Ho, Kate, and Robin S. Lee.** 2019. "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets." *American Economic Review* 109 (2): 473–522.
- Ho, Kate, and Ariel Pakes.** 2014. "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians." *American Economic Review* 104 (12): 3841–84.
- Ho, Katherine.** 2006. "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market." *Journal of Applied Econometrics* 21 (7): 1039–79.
- Ho, Katherine.** 2009. "Insurer-Provider Networks in the Medical Care Market." *American Economic Review* 99 (1): 393–430.
- Lavetti, Kurt, and Kosali Simon.** 2018. "Strategic Formulary Design in Medicare Part D Plans." *American Economic Journal: Economic Policy* 10 (3): 154–92.
- Liebman, Eli.** 2016. "Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks." Unpublished.
- Marone, Victoria R., and Adrienne Sabety.** 2021. "When Should There Be Vertical Choice in Health Insurance Markets?" Unpublished.
- Massachusetts Center for Health Information and Analysis (Massachusetts CHIA).** 2014a. *Annual Report on the Performance of the Massachusetts Health Care System*. Boston: Center for Health Information Analysis.
- Massachusetts Center for Health Information and Analysis (Massachusetts CHIA).** 2014b. *Massachusetts Hospital Profiles: Data through Fiscal Year 2012*. Boston: Center for Health Information Analysis.
- Massachusetts Health Connector.** 2013. "Commonwealth Care Program: Insurance Plans Covered Hospital Networks, 2008-2014." Public documents on hospital networks shared with the author by the Massachusetts Health Connector.
- Massachusetts Health Connector.** 2014. "Commonwealth Care Program: Administrative Enrollment and Health Insurance Claims Records." Accessed via DUA with the Massachusetts Health Connector.
- Massachusetts Health Policy Commission (HPC).** 2013. "Public Hearings – Health Care Cost Trends in the Commonwealth: Testimony from CeltiCare Health Plan." Available at <https://www.mass.gov/doc/pre-filed-testimony-response-from-celticare-0/download> (accessed 12/15/2021).
- Nevo, Aviv.** 2000. "Mergers with Differentiated Products: The Case of the Ready-To-Eat Cereal Industry." *RAND Journal of Economics* 31 (3): 395–421.
- Newhouse, Joseph P.** 2003. "Accounting for Teaching Hospitals' Higher Costs and What to Do about Them." *Health Affairs* 22 (6): 126–29.
- Newhouse, Joseph P., Mary Price, John Hsu, J. Michael McWilliams, and Thomas G. McGuire.** 2015. "How Much Favorable Selection Is Left in Medicare Advantage?" *American Journal of Health Economics* 1 (1): 1–26.

- Prager, Elena.** 2020. "Health Care Demand under Simple Prices: Evidence from Tiered Hospital Networks." *American Economic Journal: Applied Economics* 12 (4): 196–223.
- Rothschild, Michael, and Joseph Stiglitz.** 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics* 90 (4): 629–49.
- Shepard, Mark.** 2022. "Replication Data for: Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E149501V1>.
- Shorrocks, Anthony F.** 2013. "Decomposition Procedures for Distributional Analysis: A Unified Frame-Work Based on the Shapley Value." *Journal of Economic Inequality* 11 (1): 99–126.
- Song, Zirui, Dana Gelb Safran, Bruce E. Landon, Mary Beth Landrum, Yulei He, Robert E. Mechanic, Matthew P. Day, and Michael E. Chernew.** 2012. "The 'Alternative Quality Contract,' Based on a Global Budget, Lowered Medical Spending and Improved Quality." *Health Affairs* 31 (8): 1885–94.
- Starc, Amanda.** 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics* 45 (1): 198–220.
- Veiga, André, and E. Glen Weyl.** 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics* 131 (2): 1007–56.