NBER WORKING PAPER SERIES

HOSPITAL NETWORK COMPETITION AND ADVERSE SELECTION: EVIDENCE FROM THE MASSACHUSETTS HEALTH INSURANCE EXCHANGE

Mark Shepard

Working Paper 22600
http://www.nber.org/papers/w22600

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2016

I would like to thank my advisers David Cutler, Jeffrey Liebman, and Ariel Pakes for extensive comments and support in writing this paper. 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. I also thank Katherine Baicker, Amitabh Chandra, Jeffrey Clemens, Keith Ericson, Amy Finkelstein, Jon Gruber, Ben Handel, Nathan Hendren, Kate Ho, Sonia Jaffe, Tim Layton, Robin Lee, Greg Lewis, Tom McGuire, Joe Newhouse, Daria Pelech, Jim Rebitzer, Amanda Starc, Karen Stockley, Rich Sweeney, Jacob Wallace, Tom Wollmann, Ali Yurukoglu, and seminar participants at Boston College, the Boston Fed, CBO, Columbia, Chicago Booth, Duke, Harvard, NBER (Health Care), Northwestern, UPenn, Princeton, Stanford GSB, UCLA, UCSD, and Wash U St. Louis. I gratefully acknowledge data funding from Harvard's Lab for Economic Applications and Policy, and Ph.D. funding support from National Institute on Aging Grant No. T32-AG000186 (via the National Bureau of Economic Research), the Rumsfeld Foundation, and the National Science Foundation Graduate Research Fellowship. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange
Mark Shepard
NBER Working Paper No. 22600
September 2016
JEL No. I11,I13,I18,L13

ABSTRACT

Health insurers increasingly compete on their covered networks of medical providers. Using data from Massachusetts' pioneer insurance exchange, I find substantial adverse selection against plans covering the most prestigious and expensive "star" hospitals. I highlight a theoretically distinct selection channel: these plans attract consumers loyal to the star hospitals and who tend to use their high-price care when sick. Using a structural model, I show that selection creates a strong incentive to exclude star hospitals but that standard policy solutions do not improve net welfare. A key reason is the connection between selection and moral hazard in star hospital use.

Mark Shepard
John F. Kennedy School of Government
Harvard University
Mailbox 114
79 JFK Street
Cambridge, MA 02138
and NBER
mark_shepard@hks.harvard.edu

1 Introduction

Public programs increasingly use regulated markets to provide health insurance to enrollees. These types of markets now cover more than 75 million people and cost over $300 billion in U.S. programs including the Affordable Care Act (ACA), Medicare Advantage, and Medicaid managed care. These programs differ from traditional public insurance because market competition plays a key role in determining the benefits that plans cover. While competition can encourage lower costs and better quality, a longstanding concern with insurance competition is adverse selection. Adverse selection arises when individuals with (unpriced) high costs tend to select into more generous plans. This can push up the prices of generous plans, inefficiently crowding out lower-cost consumers (Akerlof 1970; Einav, Finkelstein, and Cullen 2010).[1] Selection may also may give insurers incentives to distort plan benefits to avoid unprofitable customers (Rothschild and Stiglitz 1976). Although this is a classic theoretical result, whether and how selection influences insurers' benefit design incentives is not well understood. Most the growing literature on selection studies its impact on prices given fixed contracts,[2] not on benefit competition.[3]

In this paper, I study the role of selection when insurers compete on an increasingly important benefit: their networks of covered hospitals and other medical providers. I focus on a key aspect of network quality: whether plans cover the best-regarded academic hospitals. These "star" hospitals are known as centers of advanced medical treatment and research but, partly as a result, are quite expensive (Ho 2009). By excluding them, insurers limit access to top providers but by doing so, reduce costs by steering patients to cheaper hospitals. However, insurers' incentives to balance this cost-quality tradeoff can also be influenced by selection. Understanding the role of selection is important for evaluating the recent proliferation in the ACA exchanges of "narrow network" plans, which are particularly likely to exclude academic medical centers (McKinsey 2015). But despite the importance of this issue, there is no direct evidence from past work on provider networks and adverse selection.[4]

[1] In the extreme, this can lead to unraveling of generous insurance, an outcome that past work has found to be empirically relevant (e.g., Cutler and Reber 1998; Hendren 2013, 2015; Handel, Hendel, and Whinston 2015).

[2] In addition to those already cited, papers on the pricing effects of selection in health insurance include Bundorf, Levin, and Mahoney (2012) and Hackmann, Kolstad, and Kowalski (2015). Starc (2014) and Mahoney and Weyl (2014) show how the pricing impacts differ under imperfect competition. In addition, there is a large literature testing for selection in insurance (see Cohen and Siegelman (2010) for a review).

[3] Two notable exceptions are Einav, Jenkins, and Levin (2012), who study the impact of advantageous selection on consumer credit markets; and Carey and Hall (2016), who study prescription drug formularies in Medicare Part D. In addition, there is recent theoretical work on this issue (Veiga and Weyl 2015; Azevedo and Gottlieb 2016).

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

To provide evidence, I draw upon administrative data from Massachusetts' pioneer health insurance exchange, a key model for the ACA. Using network variation across plans and within-plan over time, I show evidence of substantial adverse selection against plans covering the state's star hospitals. I show that this selection provides a strong disincentive to covering the star hospitals, even after applying the exchange's risk adjustment intended to mitigate selection incentives.

A key driver of this finding is adverse selection that occurs through a theoretically non-standard channel. Typically, economists assume that selection occurs because generous plans attract people with high *medical risk* (i.e., the sick). But in addition to medical risk, some consumers have high costs because for a given illness, they tend to *use more expensive providers*. This second dimension of costs is important because provider prices vary widely (IOM 2013; Cooper et al. 2015), and insurers typically cover the bulk of these price differences rather than passing them onto patients.[5] As a result, consumers who prefer using expensive star hospitals are more costly than (similarly sick) consumers who use less expensive alternatives. I find that plans covering star hospitals face adverse selection on both of these cost dimensions: they attract both the sick and people who tend to use star hospitals for their care.

Both of these selection channels discourage insurers from covering star hospitals and can lead to inefficiently low access to them. For instance, some consumers may value access to a star hospital in case they get seriously ill, but would be otherwise unlikely to use it. But to buy a plan covering it, they have to pool with people with much higher costs because they regularly use the star providers for their health care needs. Plans covering star hospitals differentially attract these high users, forcing them to raise prices and further crowd out infrequent users – a process that can either stabilize or lead to dropping coverage.

But distinguishing these selection channels matters for at least two reasons. First, selection on hospital preferences poses a fundamental challenge for risk adjustment, regulators' main tool for addressing selection. Risk adjustment works by measuring medical risk factors (e.g., age and diagnoses) and compensating plans that attract sicker people.[6] But even excellent risk adjustment is unlikely to offset costs arising from preferences for using star (or other expensive) providers. These preferences create residual cost variation that can lead to a breakdown of risk adjustment (Glazer and McGuire 2000).

Second, the two channels may have different cost and welfare implications. While sickness makes individuals costly in any plan, preferences for a star hospital only make enrollees costly if a plan covers that star hospital. Stated differently, preferences affect how much an individual's *costs increase* when their plan adds coverage of the star hospital. Thus, selection on provider preferences creates a form of

---

[5] The Massachusetts exchange requires that plans fully cover price differences by requiring equal copays for all in-network hospitals. However, even in less regulated settings, insurers typically cover the bulk of price variations because hospitalized patients usually exceed their plan's deductible.

[6] There is a large literature on the statistical problems of designing risk adjustment (van de Ven and Ellis 2000). However, there is limited empirical evidence on how well it works in practice. The limited evidence on selection between Medicare FFS and Medicare Advantage is mixed (see Brown et al. 2014; Newhouse et al. 2015).

selection on cost increases – analogous to the idea of "selection on moral hazard" in Einav et al. (2013).[7] My findings suggest a natural mechanism for selection on moral hazard to emerge when insurers compete on coverage of high-cost treatment options (e.g., a star hospital or an expensive drug) and patients differ in their preferences for the high-cost option relative to less expensive alternatives. This in turn matters for the efficiency cost of adverse selection. Because of the connection between selection and moral hazard, the net welfare impact of less star hospital coverage is ambiguous, and policies encouraging greater coverage may not improve welfare.

To study these issues empirically, I use data from Massachusetts' subsidized health insurance exchange, a market established in 2006 to provide coverage for low-income individuals.[8] This is a particularly nice setting to study hospital networks and adverse selection. A key barrier to past work on networks has been a lack of data from settings where plans networks differ. Most research on selection has studied employer-sponsored insurance plans that differ in cost sharing but have identical networks. The Massachusetts exchange features the reverse: all plans must set identical rules for cost sharing (and covered medical services) but can differ in their networks. This lets me study selection across plans that are nearly identical except for their provider networks.

Importantly, there is variation in coverage of the state's main star hospitals: Mass. General and Brigham & Women's hospitals, which are jointly owned by the Partners Healthcare System. Consistent with the star hospital paradigm, these hospitals are highly regarded – *U.S. News* perennially ranks them as the top two hospitals in the state – and quite expensive. To test for selection, I draw on (de-identified) enrollment and insurance claims data made available by the exchange regulator.

I start with reduced form tests for adverse selection against plans covering the star hospitals. I show that these plans attract a group with a strong attachment to the star providers: patients who have previously received outpatient care (including physician visits and other outpatient treatments) at a Partners-owned facility. Compared to other enrollees, these past Partners patients are: (1) almost five times as likely to use a star hospital for subsequent hospital admissions, (2) 28% higher cost *after* risk adjustment, and (3) 80% more likely to actively choose a plan that covers Partners. These facts suggest that Partners patients are loyal to their preferred providers and select plans partly based on their desire to continue using them. I find that this loyalty to previously used hospitals (and their affiliated physicians)

---

[7] It also relates to the findings of Bundorf, Levin, and Mahoney (2012) on selection based on individual cost increases for a PPO relative to an HMO. Einav et al. (2015) show how risk adjustment can break down when there is heterogeneity in moral hazard across individuals.

[8] This setting is distinct from Massachusetts' *unsubsidized* exchange, which has been studied by Ericson and Starc (2013, 2015a, 2015b). There has been limited past work on the subsidized exchange. Chan and Gruber (2010) study the price elasticity of consumers' insurance choices. Chandra et al. (2011; 2014) study the effects of the state's introduction of an individual mandate and of a 2008 change in cost sharing on utilization.

holds true more broadly across all hospitals in my data, suggesting that it is a general phenomenon likely to be relevant in other health insurance markets.[9]

I next study how selection played out in a case in 2012 when a large plan dropped Partners (both hospitals and affiliated physicians) from its network. This type of network change provides a natural source of evidence that has rarely been available in past research. Consistent with the selection story, I find that high-cost Partners patients were far more likely to switch plans in response to this change. Despite well-known inertia in plan choices (Handel 2013; Ericson 2014), nearly 40% of them switched plans in 2012 – compared to a switching rate of less than 5% for other enrollees. These switchers were extremely expensive (even among Partners patients), with 2011 costs 140% above the average non-switcher and risk-adjusted costs 80% higher. Most of them switched to the two remaining plans that covered Partners. This switching pattern illustrates the competitive logic of adverse selection. Dropping the star hospitals led to a large exodus from the plan, but this actually improved its bottom line (while raising costs for rivals) because the switchers were high-cost, unprofitable enrollees.

I also use the 2012 network change to test my model's key prediction of selection on moral hazard. To do so, I study cost changes for a fixed group of "stayers" in the plan after it dropped the star hospitals. I find sharp cost reductions for stayers at the start of 2012, representing a clear break from trend. These reductions are much larger for past Partners patients – about \$2,300 annually (33%) compared to \$300 (9%) for other stayers. Thus, consistent with selection on moral hazard, the types of people most likely to switch plans also experienced the largest cost reductions when they lost star hospital access.

These reduced form results provide strong evidence of the importance of adverse selection based on star hospital coverage. In the remainder of the paper, I estimate a structural model to investigate the welfare and policy implications of this selection. The model largely follows a structure developed in past work.[10] It consists of three pieces: (1) a hospital choice model, capturing patient choices under different networks, (2) a plan choice model, capturing how enrollees trade off lower plan prices (premiums) vs. better networks, and (3) a cost model, which is estimated from the claims data. Relative to past work, the main innovation is to allow for more detailed heterogeneity and use the individual-level data to capture the correlations among hospital choices, plan preferences, and costs – which are critical for adverse selection. In addition, I pay special attention to the identification of the premium coefficient in plan demand, using only within-plan variation driven by the state's subsidy rules.

---

[9] A natural question is whether this loyalty reflects state dependence (due to a cost of switching providers) or more persistent factors affecting preferences. Unfortunately, I am not able to fully separate these two channels, though other work on hospital choices suggests both are relevant (Raval and Rosenbaum 2016). Importantly, both channels have similar implications for adverse selection in the short-run, though they differ in their long-run implications for the welfare cost of limiting access to the star hospitals.

[10] See e.g., Town and Vistnes 2001; Capps, Dranove, and Satterthwaite 2003; Ho 2006, 2009; Gowrisankaran, Nevo, and Town 2015; Ho and Lee 2013.

My demand estimates imply that exchange enrollees significantly value better networks, including star hospital coverage, though by much less (in dollar terms) than has been found for higher-income consumers (Ericson and Starc 2015b). Value for star hospital coverage varies substantially and is highly skewed towards the top 10% of consumers, most of whom are past Partners patients. Consistent with the adverse selection story, these high-value consumers have both high risk-adjusted cost levels (in plans that do *not* cover Partners) and large cost increases from adding Partners coverage. Both selection channels are about equally important in explaining this group's higher costs.

Finally, I use the model to study the welfare and policy implications of the adverse selection findings. I simulate equilibrium in a simple game where insurers first choose whether or not to cover the star Partners hospitals (holding fixed other hospital coverage) and then compete on prices. The game is stylized but gives a sense of the implications for insurers' competitive incentives. The key limitation is that I do not model hospital-insurer bargaining but hold hospital prices fixed at their observed values. I also consider a robustness check in which the star hospitals exogenously reduce their prices. In simulations with policies based on the ACA, I consistently find that all or nearly all insurers drop Partners from network. Although they can increase premiums when they cover Partners, insurers lose enough money from adverse selection and higher costs for existing enrollees that profits decline.

I then analyze two simple policy changes to address selection and promote star hospital coverage. One involves modifying risk adjustment to "overpay" based on medical risk, as suggested by the risk adjustment model of Glazer and McGuire (2000). The second involves a targeted subsidy for plans that cover the star hospitals. Although the policies are different, I find quite similar results. Both changes lead to expanded coverage of Partners. However, net social welfare actually declines when this occurs. A simple analysis of consumers' value for and cost increases from Partners coverage shows the problem. Because higher-value consumers also have larger cost increases (the pattern that drives selection on moral hazard), the value curve for star hospitals lies almost entirely below the cost curve. As a result, there are few "gains from trade" in star hospital coverage for any single set of premiums sorting consumers among plans. Instead, any welfare gains would have to come from differentially sorting consumers based on their varying cost increases (as in Bundorf, Levin, and Mahoney 2012) – something the standard policies for addressing selection do not do.

The results in this paper are important for several reasons. First, they show the continued relevance of adverse selection, even in insurance markets that seek to address it through regulation and risk adjustment. They suggest a general theoretical channel – preferences for using high-cost treatment options – through which selection is likely to persist. Second, they show that limiting provider networks may be a powerful tool for insurers to avoid unprofitable customers.[11] This suggests that selection may be an important (and worrisome) driver behind the sharp rise in narrow networks in the ACA exchanges. Finally, the simulation results show the difficulty of addressing selection when it is closely linked with moral hazard. Future research should explore alternate policies that address the core sorting challenge: allocating which patients get access to expensive care from star providers.[12]

The paper proceeds as follows. Section 2 introduces the Mass. exchange setting and data. Section 3 shows reduced form evidence of adverse selection and of the specific mechanism I have highlighted (which I formalize in a simple model in Appendix A). Section 4 presents the structural model and estimates, and Section 5 shows equilibrium and counterfactual simulations. The final section concludes.

2 Massachusetts Exchange Background and Data

I study Massachusetts' subsidized health insurance exchange – called Commonwealth Care, or CommCare. Created in the state's 2006 health reform, CommCare operated from 2006-2013 to provide subsidized coverage to low-income people (below 300% of poverty) not eligible for employer insurance or other public programs.[13] Enrollees could choose among competing private plans in a centralized marketplace. Over the 2010-2013 period I focus on, the exchange averaged 170,000 enrollees – making it comparable to a large employer plan but still a small shore of the state's overall population of 6.6 million.

CommCare is a nice 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, all plans had the same patient cost-sharing rules and covered services.[14] This structure – which is more standardized than the ACA but similar to Medicaid managed care programs – lets me study plans that differ in network but are nearly identical on other dimensions.

Second, like the ACA, CommCare used sophisticated policies to counteract adverse selection. In addition to subsidies and a mandate to encourage broad participation in the market, it also employed risk adjustment based on enrollee observables.[15] Specifically, the exchange used demographics and past diagnoses to assign each enrollee a "risk score," intended to predict their relative costliness. Risk scores

[11] This finding contributes to an applied theory literature on "service-level selection" (Frank, Glazer, and McGuire 2000; Ellis and McGuire 2007), which has not previously studied networks as a selection tool.
[12] These might include demand-side policies like "tiered" copays (Prager 2016) or supply-side policies like payment incentives for physicians to steer patients more efficiently (Ho and Pakes 2014; Song et al. 2011).
[13] A separate market called "CommChoice" offered unsubsidized plans for all others. In the ACA, unsubsidized and subsidized enrollees are pooled into a single exchange, while people below 138% of poverty are eligible for Medicaid in states that have chosen to expand the program.
[14] There was an exception to this rule in two cases: (1) prescription drug formularies (for above-poverty enrollees), subject to minimum standards, and (2) a few "extra benefits" like gym memberships.
[15] CommCare also had a reinsurance program, which covered 75% of any enrollee's costs exceeding $150,000 per year. This very high cutoff meant reinsurance played a minor role, covering just 0.03% of enrollees and 1% of costs.

multiplied the plan's price (P_j), so a plan would receive $P_j \cdot \phi_i$ for someone with risk score ϕ_i. While there is debate on how well risk adjustment has worked in other settings (see Brown et al. 2014; Newhouse et al. 2015), the methods used by CommCare are state-of-the-art.

Third, Massachusetts has a clear pair of star academic hospitals: Massachusetts General Hospital (MGH) and Brigham & Women's Hospital, which are jointly owned by the Partners Healthcare System (and affiliated with Harvard Medical School).[16] *U.S. News & World Report* perennially ranks them as the top two hospitals statewide and among the top 10 nationwide. This position has given them the perception of "must-cover" hospitals that can command high prices. These high prices have been repeatedly documented for commercial insurance (Allen et al. 2008; Coakley 2013; CHIA 2014a). Table 1 shows that this pattern also holds true in CommCare. For both raw average payments per hospital admission and a severity-adjusted price measure (defined in Section 4.1), the Brigham and MGH are the two most expensive hospitals by a large margin.[17] Their per-admission prices of about \$20,000 are more than 20% higher than the next most expensive hospital and almost double the market-wide average of \$11,000.[18]

Finally, CommCare plans' hospital networks vary significantly, including in whether they cover the star hospitals. Table 2 shows coverage of the Partners hospitals by the five insurers (each of which has a single plan network). Additional statistics on the size of hospital networks are shown in Appendix Figure 1. Up to 2011, three insurers covered the star Partners hospitals. Two insurers did not, but one of these (Fallon) operated primarily in central Mass. and does not have a full Boston-area network.

My empirical work exploits a major change in Partners coverage in fiscal year 2012.[19] Spurred by an exchange policy change,[20] Network Health and CeltiCare cut their prices sharply in 2012. Although CeltiCare already had a narrow network and low cost structure (despite its covering Partners), Network Health needed to reduce costs to make this price cut feasible. To do so, Network Health dropped the Partners hospitals (and associated physicians), plus several other less prestigious hospitals.[21] Figure 1 shows that these shifts led to sharp changes in costs and hospital use patterns. After holding steady for

[16] As of 2012, Partners also owned 5 community hospitals in the Boston area and about 1,100 primary care physicians.

[17] I focus in this paper on general acute hospitals, so this list excludes specialty hospitals like Boston Children's.

[18] A natural question is whether these high prices reflect higher costs or higher markups. The answer appears to be *both*. Based on a state report of average cost per severity-adjusted admission (CHIA 2014b), the Brigham and MGH have the highest costs of any large general acute hospital, But the markups of their CommCare prices (over these average cost measures) are also the largest. Costs and markups explain about an equal share of their higher prices.

[19] CommCare's fiscal year runs from July to June, so fiscal 2012 started in July 2011.

[20] The background for the policy change is as follows. Because of federal rules, enrollees below 100% of poverty receive full premium subsidies (i.e., all plans are free). Prior to 2012, this group could choose any plan, just like higher-income enrollees. Starting in 2012, new enrollees below poverty were limited to choosing one of the two cheapest plans, which encouraged insurers to cut prices to be one of these limited choice options.

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

several years, Network Health's costs fell by 21% from 2011-2012, while costs in all other plans rose by 6%. The share of Network Health's hospital admissions going to a Partners hospital fell by two-thirds, while Partners use rose in other plans. As I show, these changes resulted from a combination of cost changes for stayers in Network Health and selection of frequent Partners users away from it.

After seeing sharply higher costs in 2012-2013, CeltiCare also dropped Partners' physicians and subsequently its hospitals in fiscal 2014, explicitly citing adverse selection as a rationale.[22] (Unfortunately, I cannot study this change because it is just at the end of my data.) Meanwhile, NHP retained Partners but had special reason to do so: Partners acquired NHP in fiscal year 2013. Thus, at the start of the ACA in 2014, only one plan covered the star hospitals and that through a vertical relationship.

Data Description: To study these issues, I use an administrative dataset on plan enrollment and insurance claims for all CommCare plans and enrollees from fiscal 2007-2013.[23] For each (de-identified) enrollee, I observe demographics, plan enrollment history, and claims for health care services while enrolled in the market. The claims include information on patient diagnoses, services provided, the identity of the provider, and the actual amounts the insurer paid for the care.

I use the raw data to construct two datasets for my analysis. The first is for hospital choices and costs. From the claims, I pull out all inpatient stays at general acute care hospitals in Massachusetts during fiscal years 2008-2013 – the period over which I have network information. I add hospital characteristics from the American Hospital Association (AHA) Annual Survey and define travel distance using the driving distance from the patient's home zip code centroid to each hospital.[24] For each hospitalization, I sum up the insurer's total payment while the patient was admitted (including both facility fees and physician professional service fees) and use this to estimate the hospital price model described in Section 4.1.

The second dataset is for insurance plan choices and total costs. I construct a dataset of available plans, plan characteristics (including premium and network), and chosen options during fiscal 2008-2013. For each enrollee x choice instance, I calculate costs over the subsequent year (from the claims data) and add on enrollee attributes, including demographics and risk scores. I consider plan choices made at two distinct times: (1) when an individual initially enrolls in CommCare or re-enrolls after a gap in coverage, and (2) at annual open enrollment when current enrollees can switch plans. A key difference between

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

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

[24] I thank Amanda Starc and Keith Ericson for sharing this data.

these two situations is their default choice. New and re-enrollees must make an active choice to receive any coverage,[25] while non-responsive current enrollees are defaulted to their current plan.

Appendix Table 1 shows summary statistics for both datasets. The data include 611,455 unique enrollees making 1,588,889 plan choices and experiencing 74,383 hospital admissions. The average age is 39.6, and about half of enrollees earn less than poverty so are fully subsidized. There is substantial flow of enrollees into an out of the market – about 11,000 people per month (or 6.5% of the market) in steady state – giving me a significant population of active choosers to assist in plan demand estimation.

3 Reduced Form Evidence of Adverse Selection

In this section, I present reduced form evidence of adverse selection against plans that cover the star hospitals in the Massachusetts exchange. I also show evidence of a key mechanism: selection driven by patients most likely to use the star hospitals, whose costs are high partly for this reason. I present three strands of evidence. Section 3.1 tests for selection by comparing *across plans* that differ in their networks. Section 3.2 then tests for selection using a large *network change* where a plan drops the star hospitals. Finally, Section 3.3 uses this same change to test for differential moral hazard across enrollees.

3.1 Star Hospital Patients and Adverse Selection

I start by testing for adverse selection by asking whether individuals with high risk-adjusted costs tend to select plans that cover Partners. I use a method similar to the positive correlation test of Chiappori & Salanie (2000), and specifically the "unused observables" approach of Finkelstein & Poterba (2014). Starting with data on plan choices, costs, and other outcomes, I run regressions of the form:

$$Y_{it} = X_{it}\alpha + Z_{it}\beta + \varepsilon_{it} \tag{1}$$

where Y_{it} are outcomes for individual *i* in year *t*, X_{it} are factors on which insurer prices can vary, and Z_{it} are other "unused" observables that insurers cannot price based on. During the 2011-13 period I analyze, the only factors in X_{it} were risk scores (used to risk-adjust payments) and income group.[26] In addition, because I run the regression across multiple years, I interact the income groups with year dummies. All standard errors are clustered at the enrollee level.

---

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

[26] Risk adjustment started in 2010, but my dataset is missing risk scores from part of 2010. Technically, insurers set a single price for all income groups, but because of subsidies, post-subsidy premiums vary across income groups. I include income groups in X_{it} to capture any effects of these varying premiums.

My goal is to use unused observables in Z_{it} that capture people's propensity to use the star hospitals. This serves both as a test of adverse selection and of the specific mechanism of selection I have described. To do so, I use a variable based on past utilization: whether an individual has previously received *outpatient care* from physicians at a star hospital or another Partners hospital (which are part of the same referral network). This measure includes both physician visits at Partners-owned practices and treatments in the outpatient wing of Partners' hospitals.[27] The measure's main limitation is that past use is only observable while patients were enrolled in the exchange. Because of this, I exclude first-time new enrollees from the analysis. Outpatient care occurs regularly enough that for the remaining sample, I observe past outpatient care use for the vast majority (87%) of people. In particular, 12% of people (and 20% of those in the Boston region) have past use at a Partners hospital.

The idea of this variable as a predictor of star hospital use is simple. When needing inpatient care, patients are likely to choose a hospital where they have an existing relationship with it or its affiliated physicians. This may be true for two reasons. First, these past relationships may *cause* patients to use the same hospital – e.g., because of physician referral patterns (see Baker et al. (2015)). Second, similar underlying factors may influence both decisions – e.g., distance, perceptions of quality. Separating these two channels (a version of the classic state-dependence vs. heterogeneity problem) is empirically challenging, and I have not attempted to do so. Importantly, both channels imply that these patients have a special affinity for the star hospitals, at least in the short run. Where they differ is in how *permanent* this preference will be if a patient loses access to the star hospitals.[28] Deciding this issue is not essential for testing for adverse selection or for analyzing market outcomes over a short horizon.

Who are these Partners patients? I show below how they differ in star hospital use, costs, and plan choices. Additionally, they are somewhat older (mean age of 42.7 versus 41.0) and sicker on observable risk score (mean of 1.29 versus 0.96, implying 33% higher predicted costs). Thus, Partners use is a correlated with observable risk but not perfectly so: these patients are represented across the full risk distribution. Interestingly, they differ little in income. They are more likely to live close to the star hospitals (median distance 11.7 vs. 35.0 miles), but the results are not completely driven by location.

Figure 2 shows the results of estimating equation (1) visually in binned scatter plot form. For each bin of risk score (x-axis), the figures show average outcomes for past Partners patients (red triangles) versus

[27] The measure is defined based on whether an individual has any outpatient claims billed to a Partners hospital prior to the time of a given plan choice. This captures visits to Partners-owned practices via the "facility fee" billed to the owning hospital. This measure includes ER visits but is essentially unchanged if these are removed – just 2% of observations are affected.

[28] State dependence implies a less durable preference, since patient-doctor relationships can be changed over time. Because of this difference, it would be interesting in future work to attempt to disentangle these two channels. Doing so would require observing shocks that exogenously change patient-physician relationships.

all others (blue circles), along with best-fit lines for each group. I also report the main coefficient (on past Partners patient status); full regression results are shown in Appendix Table 2.

The figure shows clearly that (across a wide range of observable risk) Partners patients are quite different in their hospital choices, plan choices, and costs. They are 32.2% points more likely to use one of the star hospitals (MGH or Brigham & Women's) when hospitalized, a nearly five-fold increase over the 6.6% rate for other enrollees. As a result of these hospital choices, past Partners patients' prices per admission are \$3,143 (or 29%) higher.[29] These hospital choices in turn contribute to this group having higher overall (risk-adjusted) costs – which are \$1,137 (or 28%) above the mean for other enrollees. Risk adjustment is not completely ineffective: it narrows the groups' unadjusted cost difference of \$3,286 by about two-thirds. But this still leaves a substantial gap that can lead to adverse selection.

The final piece of the unused observables test is whether these high-cost Partners patients are also more likely to choose plans that cover Partners. To limit the potential for reverse causality, I focus on "re-enrollees" who make an active plan choice upon rejoining the exchange after a break in coverage.[30] For this group, past Partners use is defined based on data from their previous coverage spell. The lower right graph in Figure 2 shows that past Partners patients are 29.8% points (or 80%) more likely to actively choose a plan covering Partners.[31] Thus, the unused observables test indicates significant adverse selection: the same group has high costs and is more likely to choose a plan covering Partners.

A natural question for these findings is whether past Partners patients' higher costs simply reflect unobserved sickness, rather than propensity to use expensive providers. While either would create adverse selection, only the latter would be evidence of the novel mechanism I have highlighted. To address this, I conduct several robustness checks. First, I run regression (1) with annual hospitalization rate (a sickness measure) as an outcome. Interestingly, after controlling for risk score, past Partners patients are *not* more likely to be hospitalized (see Appendix Table 2), suggesting little difference on unobservable sickness by this measure. Second, I run several robustness checks on the cost regressions with additional controls and for different subgroups (see Appendix Table 3). The higher costs of past Partners patients are robust to limiting the sample to enrollees with the highest-quality risk adjustment information;[32] limiting the sample to re-enrollees only; defining past use based only on physician office visits (not other forms of

[29] The results are similar if I analyze raw cost per admission instead of the severity-adjusted price measure.

[30] For current enrollees, the concern is that they select a plan covering Partners (for other reasons), then use Partners hospitals because they are available, and then remain inertial in their current plan. Focusing on re-enrollees ensure that I am at least observing active plan choices not driven by inertia. Of course, the effects could still be driven by loyalty to a plan (as opposed to loyalty to the star hospitals). I address this issue in the next section by examining plan switching choices after an insurer drops the star hospitals from network.

[31] This effect is robust to limiting the sample to re-enrollees with longer breaks from the exchange. Even among enrollees with breaks of more than two years, the effect is 21% points.

[32] Past diagnoses are unavailable for newer enrollees, so their risk adjustment is based only on age and sex. This specification limits the sample to people whose risk scores are calculated using past diagnoses.

outpatient care); and adding controls for past use of any hospital or academic hospital. These checks provide additional confidence that the effects are not simply driven by unobserved risk. Ultimately, however, it seems plausible that the group's higher costs may be driven by *both* unobserved risk and greater likelihood to use expensive providers. This conclusion would be consistent with the findings of my analysis using the structural model, discussed in Section 4.3.

3.2 Evidence from Plan Network Changes in 2012

A second way to test for adverse selection is to study plan network changes. This lets me disentangle star hospital coverage from other plan differences that could be driving the selection results just shown. They also provide more direct evidence on the selection effects of covering (or dropping) the star hospitals.

I focus on changes in 2012 that were both the largest in CommCare's history and the only time in my data period when the star hospitals were dropped from network. As discussed in Section 2, in 2012 Network Health plan dropped the Partners system (both its hospitals and doctors) as part of a strategy of shifting toward being a lower-price, more limited network plan. Other plans also changed prices at the same time but did not make significant network changes. Of course, a key assumption is that these simultaneous price changes are not driving the results, something I test below.

The network change was announced at the end of 2011, and enrollees had an opportunity to switch plans during open enrollment just before the changes took effect in 2012. I study these plan switching patterns, again using past patient status as a proxy for enrollees with a preference for the dropped hospitals. The top graph of Figure 3 plots switching rates for Network Health's enrollees during the open enrollment preceding each year. The typical switching rate is quite low (about 5%), but it spikes in 2012 to just over 10%. This entire spike is driven by past patients of the hospitals Network Health dropped. Almost 40% of past Partners patients switched away from Network Health in 2012, a more than *seven-fold* increase from adjacent years. This huge increase suggests that many patients are willing to overcome inertia and switch plans to retain access to their preferred providers.[33] Most of these switchers moved to CeltiCare and Neighborhood Health Plan, the two remaining plans covering Partners. Switching rates also spiked for past patients of the other dropped hospitals, but only to 18% (about half as much as for Partners patients). This is consistent with willingness to switch plans to retain access to a provider being a general phenomenon, but one whose effects are stronger for star hospitals.

Because the Partners patients are a high-cost group, these switching patterns had important cost implications. One way to summarize these is to compare the costs of "stayers" who remain in Network Health at open enrollment to switchers into and out of the plan. The bottom graph of Figure 3 shows

[33] One factor behind this high switching rate may be encouragement from Partners itself. Partners appears to have contacted its regular patients to let them know of the network change and encourage them to switch plans.

average costs for these three groups over the year preceding each year's open enrollment. Consistent with adverse selection, switchers out of Network Health in 2012 are dramatically more costly than stayers, with annualized costs of $8,045 compared to $3,877 for stayers. These differences are smaller but still hold for a measure of risk-adjusted costs:[34] switchers out cost $6,109 versus $3,807 for stayers. The figure makes clear that this selection pattern was not a general phenomenon across all years but unique to 2012. Comparing these risk-adjusted costs to the plan's 2011 price of $5,071 indicates that these switchers out were an unprofitable group.

Appendix Table 4 shows further statistics on stayers and switchers in 2011-2012. It makes clear that Partners patients drove the high costs among switchers out of Network Health. They represented 68% of all switchers out and had risk-adjusted costs of $6,802 in 2011 (54% above the plan average), whereas all other switchers out had risk-adjusted costs below the plan average (about $4,400). In comparison, the Partners patients who stayed with Network Health were somewhat less expensive – only $5,533 (after risk adjustment) in 2011. Thus, even among the Partners patients, dropping Partners selectively induced the highest-cost patients to switch plans.

A natural question is whether these results were driven by Network Health's (or other plans') price changes in 2012, rather than the network change. My setting provides an easy way to test this by examining selection patterns for below-poverty enrollees, for whom all plans are free (both before and after 2012). Appendix Figure 2 replicates Figure 3, with the sample limited only to below-poverty enrollees. Both switching and cost patterns for stayers/switchers are quite similar for this below-poverty sample. This suggests that the network change, not premium changes, was the driving factor.

3.3 Heterogeneity and Selection on Moral Hazard

A key prediction of my story is "selection on moral hazard": selection by the people whose costs *increase* most when their plan covers star hospital. To test this prediction, I ask whether past Partners patients – the group most likely to switch away from Network Health – also experienced the largest cost reductions when they stayed with Network Health. Of course, stayers and switchers are different people, so it is not possible to measure both cost changes and switching rates for literally the same individuals. But finding that the same characteristic predicts both outcomes would be consistent with selection on moral hazard.

I use my data's panel structure to analyze within-person cost changes from 2011-2012 for stayers in Network Health. Because the panel is unbalanced (due to churn into and out of the exchange), I use a

[34] I define risk-adjusted costs as a group's average cost divided by its average risk score. Based on the exchange's risk adjustment, a group's profitability equals $\overline{RA} \cdot P - \overline{C}$, which equals $\overline{RA} \cdot (P - \overline{C} / \overline{RA})$. Thus, $\overline{C} / \overline{RA}$ is a natural risk-adjusted cost statistic to compare to price.

regression with individual fixed effects. I restrict the sample to individuals in the market in 2011 and 2012 and who stay in a single plan during the whole period. I then run regressions of the form:

$$Cost_{i,t} = \alpha_i + \beta_t \cdot 1\{i \in NetworkHealth\} + \gamma_t \cdot 1\{i \notin NetworkHealth\} + \varepsilon_{i,t}$$

where t denotes time in bimonthly intervals. Figure 4 plots the coefficients for stayers in Network Health (β_t) and all other plans (γ_t), adding back each group's mean cost at the end of 2011. Panel A shows that costs fell sharply for stayers in Network Health at the start of 2012. The first point in 2012 is \$592 (or 14%) less than the last point in 2011 (significant at the 1% level). By contrast, costs change little for stayers in other plans, suggesting there were not important market-wide shocks at this time.[35]

Panel B shows how these cost reductions differed between Partners patients (defined as of the end of 2011) and other enrollees. Consistent with the theory, costs fall dramatically for Partners patients staying in Network Health at the start of 2012, with an estimated reduction of \$2,345 (or 33%). Cost reductions for all other enrollees are statistically significant but more modest – just \$317 (or 9%). Again, there is little evidence of sustained cost changes for either group in other plans.[36]

4 Structural Model: Demand and Costs

The reduced form evidence suggests that costly consumers who particularly like using the star hospitals select into plans covering them. In Appendix A, I present a simple model of how heterogeneity in preferences for star hospitals can lead to these results. However, to realistically study the competitive and welfare implications of this selection requires a structural model of the insurance market. I present this structural model in two steps. In this section, I present and estimate a model of hospital and insurance demand and costs, which let me specify the insurer profit function. In Section 5, I use this profit function to simulate a simple insurer competition game. Because of the richness of the data, I am able to separate these two steps, rather than relying on equilibrium assumptions for demand or cost estimation.

The model consists of three pieces: (1) a hospital choice model, (2) a plan choice model, and (3) a cost model. My selection story results from an interaction among these three factors. In this section, I show how I estimate a model that captures these interactions. Section 4.1 shows the econometric specifications and reports the estimates. Section 4.2 reports on the model's fit. Finally, Section 4.3 shows the model's implications for the correlation between preferences and costs that leads to selection.

[35] Stayers in Network Health after it dropped Partners are, of course, a selected group. In robustness checks (not shown), I have run regressions for *all individuals* (not just stayers) in Network Health vs. other plans in 2011, analogous to an intent-to-treat analysis. I find that these groups' costs are nearly identical on both levels and trends in 2011, with costs falling sharply only for Network Health enrollees at the start of 2012.

[36] The cost patterns of Partners patients in other plans alleviate the concern that this group's differential cost changes in Network Health are driven by mean reversion, since we would expect similar mean reversion in other plans.

4.1 Model Specification and Estimates

In this section, I present the model specification and estimates. I discuss the most important details, but for the sake of brevity, I relegate many details to Appendix B. The model involves three pieces: a hospital choice model, a plan choice model, and a cost model. I discuss these in turn.

Hospital Choice Model: I use the hospitalizations data to estimate a multinomial logit choice model. As noted above, I distinguish patients' hospital utility from the barriers their plan's network creates. The utility of patient *i* with diagnosis *d* for hospital *h* is:[37]

$$U_{i,d,h}^{Hosp} = \underbrace{\gamma_1\left(Z_{i,d}\right)Dist_{i,h}}_{\text{Distance}} + \underbrace{\gamma_2\left(Z_{i,d}\right)X_h + \gamma_3 PastPatient_{i,h} + \eta_h}_{\text{Hospital Characteristics x Patient Observables}} + \underbrace{\varepsilon_{i,d,h}}_{\text{Type 1 E.V. Error}} \tag{2}$$

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

$$u_{i,j,d,h}^{Hosp} = U_{i,d,h}^{Hosp} - \kappa_j \cdot 1\left\{h \notin N_j\right\} \tag{3}$$

This specification is similar to past work (e.g., Town & Vistnes 2001; Gaynor & Vogt 2003; Ho 2006). The main covariates are patient travel distance $(Dist_{i,h})$ and various hospital characteristics interacted with patient observables $(Z_{i,d})$ to allow for preference heterogeneity. These include observed hospital attributes X_h (e.g., availability of specialty services), a hospital dummy η_h capturing unobserved characteristics, and a dummy $(PastPatient_i)$ for whether patient *i* has previously received outpatient care from providers at hospital *h*.

Including past outpatient status allows me to capture relationships between patients and a hospital's physicians, which is a key source of heterogeneity in hospital choices. However, as discussed more fully in Section 3.1, this coefficient's interpretation is complicated because it picks up both state dependence and heterogeneity. To deal with this issue, I assume is that these relationships are fixed in the short run – e.g., the one-year horizon in my counterfactuals – so past use variables are held fixed in all simulations.[38] Of course, it would be nice to model the process through which these patient-provider relationships form. But doing so would introduce complicated dynamics into an already complex model. Instead, I treat these relationships as exogenous, which is sensible in the short run (but less ideal over longer horizons).

Because all covariates are observed, I estimate the model by maximum likelihood. Table 3 shows the results. Consistent with previous papers' estimates, patients dislike traveling to more distant hospitals,

[37] A full list of covariates is included in the note to Table 3. Technically, this equation is for patient *i* at a given time *t*, since many of the patient observables can vary over time.

[38] To facilitate this interpretation, the past use covariate in hospital demand is defined based only on care prior to the current plan year. I also exclude outpatient care use in the 30 days prior to admission to try to avoid picking up outpatient care directly related to the admission (e.g. pre-operative scans).

with each extra 10 miles of distance reducing a hospital's share by 29% on average. The model estimates a sizeable hassle cost for out-of-network hospitals that reduces their shares by 64% on average.[39] The table shows the largest hospital service x diagnosis interactions; the remaining coefficients are almost all significantly positive. Two sets of coefficients have implications for adverse selection. First, teaching hospitals and academic medical centers tend to attract sicker patients. Second, past outpatient use is a very strong predictor of future hospital choices. Patients choose a hospital where they have a relationship about 40% of the time, about five times higher than would be expected based on other covariates.

Plan Choice Model: I use the plan enrollments dataset to estimate a multinomial logit choice model. I treat individuals' timing of participation in the exchange as exogenous and model just their choices among exchange plans.[40] There are two times when plan choices are made: new enrollments in the exchange (including re-enrollments after a break) and a plan switching decision at annual open enrollment. For new/re-enrollee *i* choosing at time *t*, the utility of plan *j* is:

$$U_{i,j,t}^{Plan} = \underbrace{\alpha(Z_i)\cdot Prem_{i,j,t}}_{\text{Plan Premium}} + \underbrace{V\left(N_{j,t};Z_i,\beta\right)}_{\text{Network Value}} + \underbrace{\xi_{j,Reg_i,Inc_i} + \xi_{j,t,Reg_i}}_{\text{Unobs. Quality}} + \underbrace{\varepsilon_{i,j,t}^{Plan}}_{\text{Type 1 E.V. Error}} \tag{4}$$

Plan utility depends on three sets of plan attributes: premiums, networks, and unobserved quality. Premiums (which vary across enrollees as discussed below) are observed and included directly. Networks are also observed but more difficult to capture because of their high dimensionality. To model their role, I include two sets of terms in $V(.)$. First, I follow the literature by including an expected "network utility" measure from the hospital choice model into plan demand.[41] Second, I include a variable for whether the plan covers the hospitals with which the consumer has past outpatient relationships (or the share covered if there are multiple). In interpret this variable as picking up the utility of access to a hospital's physicians for outpatient care, though it may also pick up misspecification in the calculation of network utility.

The final covariates are plan dummy variables capturing unobserved quality. I use these to assist in identification of the premium coefficient. The classic concern is that prices are correlated with unobserved quality due to strategic firm pricing. Most papers address this problem with instruments, but I take a different approach. I use *within-plan* variation induced by CommCare's subsidies. The key fact is that subsidies make all plans free for below-poverty enrollees, while prices differences pass through to higher-income enrollees. This structure also creates differential premium *changes* over time, and I use

[39] A 63% reduction from being out of network may seem low. However, it is consistent with a basic fact in the data: 25% of hospital options are out of network, and these are chosen 8% of the time (about 1/3 of 25%).
[40] This assumption is reasonable because eligibility is determined by exogenous factors (e.g., income and job status), and generous subsidies encourage participation by the eligible. Further, the main variable likely to affect participation in the exchange, the premium of the cheapest plan, is set directly by the exchange's subsidy rules.
[41] This method was developed by Town and Vistnes (2001) and Capps, Dranove, and Satterthwaite (2003). See Appendix B for a formal definition of the network utility measure.

these changes for identification. Econometrically, I include plan-region-year dummies to absorb premium differences arising from plan pricing. I also include plan-region-income group dummies to absorb persistent demand differences across incomes. The remaining variation is from within-plan differential premium changes across income groups caused by the subsidy policy.

This identification strategy is analogous to difference-in-differences (in a non-linear model). Thus, the key assumption is parallel trends in demand across income groups. Figure 5 shows a test of parallel trends. It plots monthly choice shares for new enrollees around price changes (at time 0), separately for plans that cut prices (top graph) and increase prices (bottom graph). Consistent with the key assumption, market shares are flat and parallel for both groups at all times except time 0. At this time, demand shifts in the expected direction for premium-paying enrollees but is unchanged for below-poverty enrollees.

The model so far applies to new/re-enrollees making active plan choices. I make one adjustment for current enrollees, whose plan switching choices may be affected by inertia (Handel 2013; Ericson 2014). In addition to the utility in (4), I include a dummy variable for their current plan to capture inertia in a simple way. I allow the coefficients to vary with enrollee observables. This setup ensures that the model will match average switching rates, but the coefficients themselves may pick up both true inertia and persistent unobserved heterogeneity. For my purposes, it is not clear that is important to distinguish these factors. Doing so would matter primarily for dynamic price competition, which I do not model.

I estimate the model with GMM, using moments similar to those in Berry, Levinsohn, and Pakes (2004). See Appendix B for the moment formulas. Table 4 shows the results. Premiums (in dollars per month) enter negatively and significantly for all groups. (I normalize the average premium coefficient to -1.0 so remaining coefficients can be interpreted as dollar values.) Overall, enrollees are quite premium sensitive: a $1 increase decreases an average plan's market share by 3.0%. Enrollees also significantly value better hospital networks, both through the network utility and past-used hospital variables. Network utility is normalized so 1.0 is the average utility loss for Boston-area enrollees when Network Health dropped Partners in 2012. Thus, a typical Boston-area enrollee with no relationship with Partners hospitals would value it a fairly modest $6-8 per month. However, the past use variables mean that people with existing relationships at a Partners hospitals value it much more highly.

As expected, there is substantial inertia in consumers' plan switching decisions, with an estimated switching cost of $96.8 (or equivalently, a 94.6% probability of being passive).[42] Though large, this estimate is actually a bit smaller than Handel's (2013) estimate of $2,032 per year (or $169 per month). Interestingly for selection on networks, the estimated inertia decreases sharply when a plan drops an enrollee's past used hospital from network. This is consistent with the finding of Ho, Hogan, and Scott-Morton (2015) that enrollees pay more attention when their plan makes major changes.

[42] Appendix B shows the assumptions under which this probability of being passive is computed.

Cost Model: The final piece of the structural model is costs. The goal is to capture how individual-level costs vary across different plans with varying hospital networks. I take two different approaches for inpatient hospital costs versus all other ("non-hospital") costs. For hospital costs, I start by conditioning on each person's set of observed hospitalizations (and associated diagnoses). I then use the hospital choice model to predict how hospital choices (and therefore costs) would change under different plan networks. This method has the advantage of letting me capture the correlation between hospital use and enrollee attributes (which determine selection) in a rich, nonparametric way. Nonetheless, this approach assumes networks do not affect the *number* of hospitalizations, only the hospitals chosen when sick.[43]

I first estimate a model of insurer-hospital prices for inpatient care using the payment data. Because actual payment rules are unknown (and likely quite complicated), I follow past work in estimating a model of *average* prices for each insurer-hospital-year, controlling for patient severity. The details of the specification and method are in Appendix B. The results are hospital prices, $P_{j,h,t}$, and a patient severity metric, $\omega_{i,t,a}$, that captures the relative costliness of each admission in the data.

With these estimates, I define the hospital costs of enrollee *i* in plan *j* (with network N_j) at time *t* as:

$$c_{i,j,t}^{Hosp}\left(N_j\right) = \sum_{a=1}^{NAdmits_{i,t}} \omega_{i,t,a} \cdot \left(\sum_h P_{j,h,t} \cdot s_{i,j,d,h}\left(N_j\right) \right) \tag{5}$$

where $s_{i,j,d,h}\left(N_j\right)$ is the implied probability for hospital *h* from the hospital choice model. Notice that this model lets me calculate costs both in alternate plans and alternate networks for a given plan.

For non-hospital costs, I do not have a model of provider choices and prices through which to define costs analogously. Instead, I take a simpler approach. I estimate plan effects from a Poisson regression of non-hospital costs on individuals' diagnoses, demographics, and dummies for each plan-region-year (see Appendix B for specification details). Let $\nu_{j,Reg,t}$ be the estimated (multiplicative) plan effects from this regression. If an individual enrolls in plan *k* instead of plan *j*, I assume that their observed non-hospital costs are scaled by a factor $\left(\nu_{k,Reg,t} / \nu_{j,Reg,t}\right)$.

For my equilibrium simulations, I need one additional adjustment to capture how non-hospital costs change if a plan adds (or removes) Partners from its network. Non-hospital costs may change if patients substitute to Partners physicians for outpatient care.[44] To estimate this effect, I assume that non-hospital costs change in proportion to the *average* percent change in hospital costs for similar enrollees when

[43] This assumption is likely a good first approximation but is not perfect. Recent evidence from Gruber and McKnight (2014) finds small reductions in the number of hospitalizations in limited network plans. If applicable in my setting, my model will somewhat understate the cost savings from plans' limiting their networks.
[44] Past structural work on hospital networks has generally either ignored non-inpatient costs or assumed that they did not change with the hospital network. My method, though quite rough, improves on this baseline.

Partners is added/dropped (which I can compute from the hospital cost model above). I compute this average separately by plan, year, region, and past Partners patient status, to allow the effect to differ for these groups. I also scale down the effect by a factor capturing the typical relationship between hospital and non-hospital costs.

Putting these together, total costs, $c^{Tot}_{i,j,t}\left(N_j\right)$, equal the sum of hospital costs, non-hospital costs, and a measure of plans' variable administrative costs, which I estimate from plan financial reports.[45]

Plan Profit Function: Putting the three pieces above together, I have everything needed to specify the plan's profit function. Profits as a function of plan price $(P_{j,t})$ and network $(N_{j,t})$ equal:

$$\pi_{j,t}\left(P_{j,t}, N_{j,t}\right) = \sum_i \left(\phi_{i,t} P_{j,t} - c^{Tot}_{i,j,t}\left(N_{j,t}\right)\right) \cdot D_{i,j,t}\left(\mathbf{Prem}(\mathbf{P}), \mathbf{N}\right) \tag{6}$$

where $\phi_{i,t}$ is the enrollee's risk score and $D_{i,j,t}\left(\mathbf{Prem}(\mathbf{P}), \mathbf{N}\right)$ is enrollee demand for plan *j* at time *t*. Demand is in units of member months and equals the product of the enrollee's duration in the exchange (treated as exogenous) times the choice probability implied by the plan choice model. Demand is a function of all plans' premiums and networks (denoted by bold). Premiums are in turn set based on plan prices, applying a given subsidy rule.

4.2 Model Fit

Appendix Figure 3 shows the model fit for plans' average monthly medical costs per enrollee. Both error in the model's demand and cost functions may cause average costs to differ from the data. Nonetheless, the fit is quite good, with an R^2 at a plan-year level of 0.926. Importantly, the model captures very well the large fall in costs for Network Health in 2012 when it dropped Partners. The largest errors are predicting too high costs for CeltiCare in 2010 and 2011 when it was a new plan and had very low enrollment. However, the model does capture its large cost increase in 2012 after Network Health dropped Partners and CeltiCare (which still covered Partners) had an influx of high-cost enrollees.

Appendix Figures 4-6 show how the model matches more detailed cost and demand patterns for Network Health around its 2012 dropping of the Partners hospitals. This is not an out-of-sample fit (since the model is estimated on all years) but tests whether the model can capture the heterogeneity in outcomes for different groups that drives selection. The fit is credibly good for plan switching rates, cost changes for stayers vs. other enrollees, hospital choice patterns, and costs per hospital admission. The model's ability to capture these patterns around the 2012 network change adds confidence to its use for simulating counterfactual demand and cost outcomes.

[45] I estimate a regression of reported administrative costs on a plan's total enrollment. I find an almost perfect linear fit with a coefficient of about $30 per member-month. I take this as the variable administrative cost for all plans.

4.3 Analysis of Adverse Selection

I now use the model estimates to analyze the key relationship driving adverse selection: the correlation between costs and demand for the Partners hospitals. (I include both the star and other Partners hospitals because these are typically covered or excluded in a bundle.) To do so, I define consumer willingness-to-pay (WTP) metric for Partners coverage as the utility of adding Partners to a network (based on $V\left(N_{jt};\, Z_i, \beta\right)$ in the plan utility function in (4)) divided by marginal utility of money (based on the price coefficient). This WTP varies across plans and years based on the other hospitals they cover. For simplicity, I focus on WTP calculated for a single plan (Network Health) in 2012.[46]

Table 5 shows WTP and cost statistics for consumers sorted by increasing WTP for Partners coverage. About 85% of enrollees have relatively little value for Partners coverage (less than \$10 per month).[47] But value for Partners rises sharply in the top 10% of enrollees, with the top 2% valuing Partners at \$52.8 per month. For these enrollees – most of whom are past Partners patients – Partners coverage plays a critical role in their plan choice.

The remainder of Table 5 shows how these differences in value for Partners correlate with costs. Columns (2) and (3) show insurer costs for each group when the plan does *not* cover Partners. Both raw and risk-adjusted costs generally rise with WTP. For instance, the top 2% of consumers by WTP have risk-adjusted costs of \$356.6 per month, about \$52 (or 17%) higher than those with below-median WTP. The gradient in unadjusted costs is even steeper. These results indicate that consumers who value Partners are costly even without Partners in network – consistent with what I call "selection on cost level."

Columns (4)-(5) show that WTP also correlates with cost increases (or moral hazard) when the plan adds Partners to network – consistent with selection on moral hazard. The highest-WTP group has cost increases of \$63.9, which is \$55.5 (or more than *six times*) larger than the cost increase for the below-median WTP group. These figures indicate the close link between WTP for and cost increases from Partners. This occurs because both WTP and cost increases are driven by enrollees' propensity to use Partners. Thus, selection on moral hazard is a natural outcome in this setting. These estimates suggest that selection on cost levels and moral hazard are about equally important in explaining the higher risk-adjusted costs of high-WTP enrollees in plans that cover Partners (both explain about \$50).

A final insight from Table 5 is that WTP for Partners coverage (column (1)) falls short of the insurer cost increase (column (4)) across the entire distribution of consumers. This indicates that there are very few gains from trade in covering Partners. Selection on moral hazard (the correlation between WTP and

[46] Results for other plans are qualitatively similar. I also exclude below-poverty enrollees from this calculation because I cannot estimate their marginal utility of money (since they do not pay premiums).

[47] This is partly because of enrollees living outside of the Boston area, but even among those within 30 miles of Boston, 75% of them have WTP of \$10 or less.

ΔCost) plays a key role in making this true. If everyone had the average ΔCost of $18.4, then there would be gains from trade for the top 10% of consumers by WTP. Instead, these consumers also have differentially high ΔCost, which exceed even their high WTP.

One reason for low gains from trade between insurers and consumers is that the Partners hospitals extract some surplus in their price markups. To account for these markups, I draw on a state estimate of hospitals' average costs per-admission (CHIA 2014b).[48] Column (6) shows the net social cost increase, after subtracting the increase in Partners net revenue.[49] This lowers the ΔCost for the highest-WTP groups, making it socially optimal for them to get Partners coverage. Thus, the star hospitals' markups may be an additional rationale (beyond adverse selection) for encouraging their coverage.

5 Equilibrium and Counterfactual Policies

This section uses the demand and cost estimates to simulate equilibrium in a model of insurance competition. I use this to examine the impact of different policies used to address adverse selection in insurance exchanges. In general, insurer competition on prices and hospital networks may be extremely complicated. To make progress, I focus on a static model where insurers compete only on price and coverage of the star Partners hospitals, holding fixed hospital-insurer prices and other aspects of the network. Although stylized, this model goes beyond most past empirical work on selection, which studies pricing holding fixed product characteristics.

5.1 Equilibrium Simulations: Method and Results

Consider a model of insurance market equilibrium for a particular year (e.g., 2012) in the exchange. As in Massachusetts, I assume that each insurer offers a single plan with exchange-specified consumer cost sharing and covered service rules. I condition on each plan's past history, including past network coverage and the set of current enrollees entering the year. I also hold fixed (at observed values) each plan's network and payment rates for all non-Partners hospitals. Before the year, the exchange announces policies (e.g., subsidy and risk adjustment rules). Insurers then compete in the following game:

Insurer Competition: 1. Insurers choose whether to cover Partners hospitals

2. Insurers set plan prices

[48] The measure is of hospitals' "inpatient cost per case mix adjusted discharge." The calculation, which is based on hospitals' cost reports to the state, is intended to be a comprehensive measure of average hospital costs (including some fixed costs), excluding physician compensation and graduate medical education costs. Based on these estimates, the cost per admission at the two star Partners hospitals in 2012 were about $12,500 (MGH) and $13,800 (Brigham), implying margins of about 30-35% relative to my estimated prices.

[49] This calculation is imperfect because it does not account for reductions in net revenue for non-Partners hospitals. The latter, however, are likely to be small; non-star hospitals often have low or negative margins (Ho 2009).

Consumer Demand: 3. Consumers choose plans

4. Sick consumers choose providers (based on plan network)

I assume that insurers observe networks from stage 1 when setting prices and that they have full information on all demand and cost functions.

Insurers make choices to maximize profits, using the profit function specified in Section 4. However, there is one additional simulation issue: how to incorporate the dynamic effects of enrollee inertia. When a plan lowers its price and attracts more enrollees today, it increases its future demand because some enrollees will passively stick with the plan in later years. This can lead to an invest-then-harvest equilibrium in which plans cycle between low and high prices. I choose not to specify a fully dynamic model, which would be both complicated and unrealistic unless it modeled uncertainty about policy changes (which occurred frequently in Massachusetts). Instead, I take a simple static approach that approximates the dynamic incentives involved. I assume that enrolling someone today increases future profits in proportion to the person's future duration in the market and the current profit margin on the individual. This "future profit effect" gives insurers an added incentive to keep prices low and helps offset the lower price elasticity of demand due to inertia. Appendix C shows the modified pricing first-order conditions and lays out additional details of the simulation method.[50]

In Nash equilibrium, each insurer sets prices in step 2 to satisfy its FOC given all other insurers' prices and networks. In step 1, they choose Partners coverage knowing the pricing equilibrium that will prevail for each network. For Partners coverage, I assume a binary choice: either sticking with their actual coverage of Partners or adding/dropping all of the Partners hospitals. I do not model the vertical relationship between Partners and Neighborhood Health Plan (NHP) but allow it to flexibly cover/drop Partners.[51] Nash equilibrium occurs at a set of networks if no insurer wishes to unilaterally deviate. While uniqueness is not guaranteed, I typically find a single equilibrium. When I do not, I report all equilibria.

Table 6 shows equilibrium insurer choices for several simulations. The top panel shows equilibrium under the actual Massachusetts policies in 2011, comparing these to the observed outcomes. While I would like to perform a similar test for other years, data limitations and policy complexities make this challenging.[52] The model's prices match extremely well. But this occurs largely because the exchange required plans to price within a narrow range, and all plans bid near the range's min or max. Nonetheless, the model captures well which insurers priced near the min vs. the max. For networks, the model predicts just one plan (CeltiCare) willing to cover Partners, while in reality Network Health and NHP also covered

[50] An alternate assumption would be to ignore inertia and treat all consumers as active choosers. This would make the model static but would likely overstate the price elasticity of demand.

[51] I also hold fixed the observed choice of Fallon (which is unavailable in most of Boston) not to cover Partners.

[52] Prior to 2010, the pricing process was more complicated and involved some negotiation with the exchange. In 2010, I am missing data on risk scores. And in 2012-13, the exchange introduced a limited choice policy that creates auction-like dynamics that would be much more complicated to model.

Partners in 2011. However, Network Health did drop Partners in 2012, and Partners was in talks to acquire NHP at the time, a factor I do not model. Interestingly, the model rationalizes CeltiCare's surprising decision (as the low-price plan) to cover Partners. CeltiCare is willing to do so because of the binding price floor. Without a price floor, CeltiCare instead cuts its price and drops Partners.

Because many of Massachusetts' policies do not continue under the ACA, I perform the rest of the analysis using rules closer to the ACA's. Specifically, I include only the 100-300% poverty population, set subsidies as a flat amount for all plans, and do not impose min or max prices.[53] Panel B of Table 6 shows the results. Under ACA-like policies, all plans drop Partners, and this result is robust across all the simulation years, 2011-2013. This occurs despite the fact that insurers can flexibly set plan prices. If an insurer deviates to cover Partners, its costs go up, and it tends to attract enrollees with high risk-adjusted costs. If the plan raises its price to compensate, it reduces demand among a large number of lower-cost enrollees. The net result is that plans lose money by covering Partners.

A limitation is that these simulations hold fixed Partners' hospital prices. This may be reasonable for the small CommCare market (covering about 3% of the state population), and indeed, I find that Partners did not lower its prices much after Network Health dropped it in 2012. However, if plans in a broader array of markets dropped Partners, it might be forced to respond. While a full model of hospital-insurer bargaining is beyond the scope of this paper, I can consider the robustness of my results to alternate (exogenously set) Partners prices. Panel C shows what happens if Partners hospitals completely eliminate their markups over a measure of their average per-patient costs (see definition in Section 4.3). This represents an average price cut of about 30%. For 2011 and 2012, Partners coverage still fully unravels. In 2013, only NHP covers Partners but is nearly indifferent with dropping it (it would do so if Partners prices were slightly higher). Thus, the limited coverage of Partners is robust to sizeable price cuts.

5.2 Social Welfare Function

To analyze the welfare implications of different equilibria, I need a social welfare function. My starting point is a social surplus approach, in which welfare equals consumer surplus plus insurer and Partners profits, minus government costs. This assumes Partners' profits are valued by society dollar-for-dollar, though I can consider other assumptions. I make one key adjustment in the calculation of consumer welfare. I choose to exclude the switching cost from consumers' plan utility, treating it as welfare-irrelevant inattention. Recall that the estimated switching costs were lower when a plan dropped a

[53] The remaining differences with the ACA are the lack of unsubsidized enrollees (who are about 20% of ACA enrollees) and the absence of multiple generosity tiers. Because of these differences, the simulations should be seen as illustrative of the economic forces involved, not a perfect prediction of outcomes in the ACA.

consumer's hospital, and I do not want the welfare analysis to be driven by this difference.[54] Once I exclude switching costs, however, the standard inclusive value formula for expected utility in a logit model does not apply. Instead, I define expected plan value for consumer *i* as:

$$ConsValue_i = \tfrac{1}{-\alpha_i} \sum_j \hat{s}_{ij}^{Plan} \cdot \hat{U}_{ij} \tag{7}$$

where α_i is the individual's premium coefficient, $\hat{s}_{ij}^{Plan}$ is the model's predicted share for consumer *i* choosing plan *j*, and $\hat{U}_{ij}$ is plan utility excluding switching costs and the logit error.

5.3 Policy Counterfactuals

I next simulate policy changes to address the unravelling of Partners coverage. I focus on two policies that are standard responses to adverse selection. The first policy builds on an idea of Glazer and McGuire (2000) to modify risk adjustment by "over-paying" for high-risk types (and paying less for low-risk types). The logic is that by over-paying for observable risk, the exchange can compensate for adverse selection on unobserved costs. The key assumption is that the same plans (here, plans covering Partners) face adverse selection on *both* observed and unobserved costs – something that appears to be true in this setting. To implement this, I multiply all risk scores above the mean by a factor $(1+\phi)$, divide below-mean risk scores by the same factor, and renormalize the distribution to have mean 1.0.

This modified risk adjustment has the advantage of being generally applicable: it can mitigate adverse selection on benefit generosity even if regulators are unsure which benefits lead to the selection. Alternatively, I consider a second, more targeted policy: an extra subsidy for plans that cover the Partners hospitals. The logic of this policy as a way to encourage Partners coverage is straightforward. Of course, regulators may be uncomfortable with such an explicit intervention in favor of certain hospitals.

Table 7 shows the equilibrium outcomes for modified risk adjustment (top table) and subsidies (bottom table), as well as social welfare and its components.[55] All of the simulations build on the ACA population and policies (Panel B of Table 6), and for simplicity I show results just for 2012. Two main results emerge. First, these policies can reverse the unravelling of Partners coverage. The over-payment factor ϕ needs to be fairly large (about 100%) to do so, but even a modest \$4 monthly subsidy gets one plan to cover Partners (and all plans cover it with a \$20 subsidy). This is a fairly small amount compared to the exchange's existing subsidies of over \$300 per month. Thus, as a technical matter, policymakers can make simple changes to address narrower networks arising from adverse selection.

[54] I have also done the welfare analysis with switching costs included. The results are qualitatively similar, but past Partners' patients value of coverage is higher because of the switching cost interaction.
[55] The level of consumer surplus is unknown (it depends on their value of participating in the exchange, which I cannot observe), so I report values differenced relative to the baseline policy.

However, the second result is that social surplus uniformly declines under these modified policies. Although consumers, insurers, and Partners all benefit, a sharp increase in government costs more than offsets these gains. Despite the different natures of the two policies, the results are qualitatively similar. One difference is that insurer profits (and government costs) increase somewhat more with modified risk adjustment. This is consistent with the findings of Starc (2014) and Mahoney & Weyl (2014) that adverse selection lowers insurer markups. Flattening the cost curve by strengthening risk adjustment may thus lead to higher markups, which raises government costs because subsidies are linked to prices.[56]

Underlying these results are the small estimated gains from trade in covering Partners, as shown in Section 4.3. While a subset of enrollees have high willingness-to-pay (WTP) for Partners coverage, these enrollees also tend to have high cost increases (ΔCost) from it because they use the star hospitals frequently. The gains from trade (= WTP minus ΔCost) are negative for most people and positive only for the highest-WTP groups after accounting for the Partners markup. Importantly, these are averages for consumers with a given WTP for Partners access. Within a single value of WTP, some consumers may have smaller ΔCost and therefore positive net value of Partners coverage. However, a single plan price can only sort consumers based on WTP, not on ΔCost.[57] In settings where ΔCost varies, differential plan prices for different groups may be needed to improve welfare (Bundorf, Levin, and Mahoney 2012).

An important caveat is that if Partners' net revenue has positive externalities (because it funds socially valuable teaching and research), the welfare effects could be different. Some simple arithmetic shows that each $1 of Partners profits would need to produce a positive externality of at least $0.90 for social surplus to improve in any of the counterfactuals shown. Whether such a large externality is plausible is difficult to judge but is an interesting question for future research.

6 Conclusion

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

[56] This occurs because of the "price-linked" subsidy design used by the Massachusetts and ACA exchanges (Jaffe and Shepard 2016; Tebaldi 2016).

[57] Indeed, sorting is likely to be even less optimal than based on WTP for Partners. Plans are bundles of many attributes and consumers sort based on their overall valuation of the bundle. Some consumers with low WTP for Partners may nonetheless choose a plan based on its other attributes. But once they have free access to the star hospitals, they may decide to use them, driving up cost at the insurer's expense.

hospitals. And these consumers have high costs precisely *because* they use the expensive star providers for care. This creates selection on a dimension of costs unlikely to be offset by medical risk adjustment.

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

The findings also have general implications for how economists think about adverse selection in health insurance markets. My results suggest that consumer preferences for high-cost treatment options – star hospitals in my study, but the same idea could apply to any expensive provider, drug, or treatment – can naturally lead to adverse selection, and specifically selection on moral hazard. Selection on moral hazard is not just an empirical curiosity but affects welfare and policy implications. Typically, economists think of adverse selection as leading to too little access to (or enrollment in) generous insurance, creating a rationale for mandates or subsidies. But selection on moral hazard complicates the analysis because people with the greatest demand for a generous benefit also have the largest cost increases from it. As a result, subsidies for generous coverage may not improve welfare. My model simulations illustrate this possibility. Although covering the star Partners hospitals leads to adverse selection – so much so, that no plan covers Partners at baseline – policies that subsidize coverage of it actually reduce social surplus.

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

References

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

Allen, Scott et al. 2008. "A Healthcare System Badly out of Balance." *The Boston Globe*.

Azevedo, Eduardo M, and Daniel Gottlieb. 2016. "Perfect Competition in Markets with Adverse Selection." Working Paper.

Baker, Laurence C, M Kate Bundorf, and Daniel P Kessler. 2015. "The Effect of Hospital/Physician Integration on Hospital Choice." *NBER Working Paper*.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112(1): 68–105.

Breyer, Friedrich, M. Kate Bundorf, and Mark V. Pauly. 2011. "Health Care Spending Risk, Health Insurance, and Payment to Health Plans." *Handbook of Health Economics* Vol. 2.

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." *The Rand Journal of Economics* 34(4): 737–63.

Carey, Colleen, and Martha Van Rensselaer Hall. 2016. "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D." *American Economic Journal: Economic Policy*.

Chan, David, and Jonathan Gruber. 2010. "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review, Papers & Proceedings* 100(2): 292–96.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2011. "The Importance of the Individual Mandate -- Evidence from Massachusetts." *New England Journal of Medicine* 364: 293–95.

———. 2014. "The Impact of Patient Cost-Sharing on Low-Income Populations: Evidence from Massachusetts." *Journal of Health Economics* 33(1): 57–66.

(CHIA) Massachusetts Center for Health Information & Analysis. 2014a. "Annual Report on the Performance of the Massachusetts Health Care System."

———. 2014b. "Massachusetts Hospital Profiles: Data through Fiscal Year 2012." See Figure 5 (p. 13) with numbers in data appendix available via request to CHIA.

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

Coakley, Martha. 2013. "Examination of Health Care Cost Trends and Cost Drivers." Available at http://www.mass.gov/ago/docs/healthcare/2013-hcctd.pdf.

Cohen, Alma, and Peter Siegelman. 2010. "Testing for Adverse Selection in Insurance Markets." *Journal of Risk and Insurance* 77(1): 39–84.

Cooper, Zack, Stuart Craig, Martin Gaynor, and John Van Reenen. 2015. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." *NBER Working Paper*.

Cutler, David M., and Sarah J. Reber. 1998. "Paying For Health Insurance: The Trade-off Between Competition and Adverse Selection." *The Quarterly Journal of Economics* 113(2): 433–66.

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, 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. 2015. "Beyond Statistics : The Economic Content of Risk Scores." *NBER Working Paper*.

Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. "Contract Pricing in Consumer Credit Markets."

Econometrica 80(4): 1387–1432.
Ellis, Randall P., and Thomas G. McGuire. 2007. "Predictability and Predictiveness in Health Care Spending." *Journal of Health Economics* 26(1): 25–48.
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. 2013. "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange." *NBER Working Paper*.
———. 2015a. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics*.
———. 2015b. "Measuring Consumer Valuation of Limited Provider Networks." *American Economic Review, Papers & Proceedings* 105(5): 115–19.
Finkelstein, Amy, and James M. Poterba. 2014. "Testing for Asymmetric Information Using 'Unused Observables' in Insurance Markets: Evidence from the U.K. Annuity Market." *Journal of Risk and Insurance* 81(4): 709–34.
Frank, Richard G., Jacob Glazer, and Thomas G. McGuire. 2000. "Measuring Adverse Selection in Managed Health Care." *Journal of Health Economics* 19: 829–54.
Gaynor, Martin, and William B Vogt. 2003. "Competition Among Hospitals." *Rand Journal of Economics* 34(4): 764–85.
Glazer, Jacob, and Thomas G McGuire. 2000. "Optimal Risk Adjustments in Markets with Adverse Selection: An Application to Managed Care." *American Economic Review* 90(4): 1055–71.
Glied, Sherry. 2000. "Managed Care." *Handbook of Health Economics* Vol. 1: 707–53.
Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry." *American Economic Review* 105(1): 172–203.
Gruber, Jonathan, and R McKnight. 2014. "Controlling Health Care Costs Through Limited Network Insurance Plans: Evidence from Massachusetts State Employees." *NBER Working Paper*.
Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski. 2015. "Adverse Selection and an Individual Mandate : When Theory Meets Practice." *American Economic Review* 105(3): 1030–66.
Handel, Benjamin. 2013. "Adverse Selection and Inertia in Health Insurance Markets :" *American Economic Review* 103(7): 2643–82.
Handel, Benjamin, Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges : Adverse Selection versus Reclassification Risk." *Econometrica* 83(4): 1261–1313.
Hendren, Nathaniel. 2013. "Private Information and Insurance Rejections." *Econometrica* 81(5): 1713–62.
———. 2015. "Knowledge of Future Job Loss and Implications for Unemployment Insurance." Working Paper.
Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2015. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." *NBER Working Paper*.
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 U.S. Medical Care Market." *Journal of Applied Econometrics* 21(November): 1039–79.
———. 2009. "Insurer-Provider Networks in the Medical Care Market." *American Economic Review* 99(1): 393–430.
Ho, Katherine, and Robin S Lee. 2013. "Insurer Competition and Negotiated Hospital Prices." *NBER Working Paper*.
(HPC) Massachusetts Health Policy Commission. 2013. "Health Care Cost Trends in the Commonwealth, Hearing, September 16, 2013." Available at http://www.mass.gov/anf/docs/hpc/celticare.pdf.
(IOM) Institute of Medicine. 2013. *Variation in Health Care Spending: Target Decision Making, Not Geography*.
Jaffe, Sonia and Mark Shepard. 2016. "Price-Linked Subsidies and Health Insurance Markups." Working Paper.

Mahoney, Neale, and E Glen Weyl. 2014. "Imperfect Competition in Seleciton Markets." *NBER Working Paper*.
McKinsey. 2015. "Hospital Networks : Evolution of the Configurations on the 2015 Exchanges."
Newhouse, Joseph P, Mary Price, 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. 2016. "Tiered Hospital Networks, Health Care Demand, and Prices." Working Paper.
Raval, Devesh, and Ted Rosenbaum. 2016. "Why Are Distance and Switching Costs Important? Evidence from Childbirth." Working Paper.
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.
Song, Zirui, Dana Gelb Safran, Bruce E. Landon, Yulei He, Randall P. Ellis, Robert E. Mechanic, Matthew P. Day, and Michael E. Chernew. 2011. "Health Care Spending and Quality in Year 1 of the Alternative Quality Contract." *The New England Journal of Medicine* 365(10): 909–18.
Starc, Amanda. 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics* 45(1): 198–220.
Tebaldi, Pietro. 2016. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." Working Paper.
Town, Robert, and Gregory Vistnes. 2001. "Hospital Competition in HMO Networks." *Journal of Health Economics* 20(5): 733–53.
Veiga, André, and E Glen Weyl. 2015. "Product Design in Selection Markets." Working Paper.
Van de Ven, Wynand P M M, and Randall P Ellis. 2000. "Risk Adjustment in Competitive Health Plan Markets." *Handbook of Health Economics* Vol. 1.

Figure 1. Costs and Hospital Use around 2012 Network Health Changes

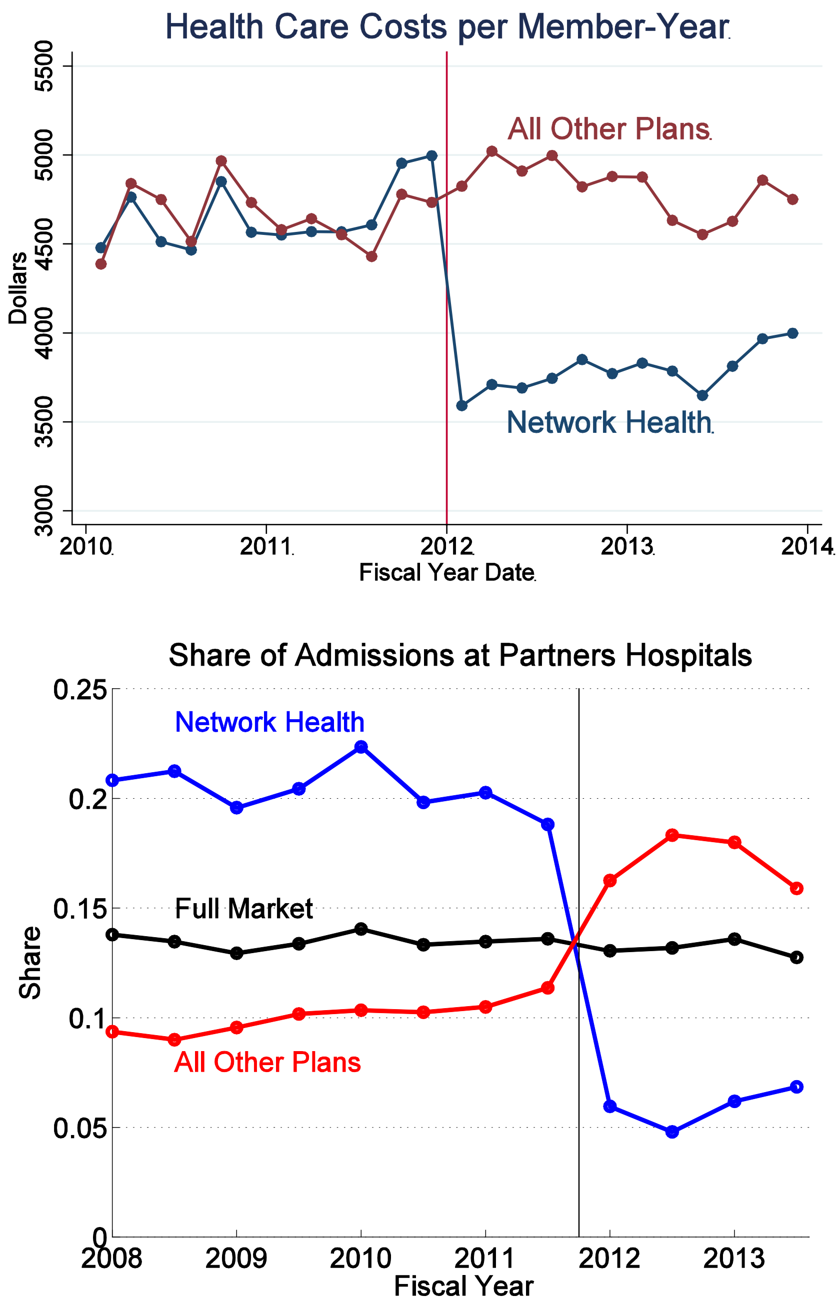

NOTE: These graphs show summary statistics for costs and Partners hospital use in Network Health – which drops Partners in 2012 – and all other plans. The top graph shows average costs for the enrollees in each plan / group of plans. Costs fall sharply for Network Health at the start of 2012. The bottom graph shows the share of plans' hospital admissions that occur at a Partners hospital (which includes the two star hospitals and several affiliated community hospitals). This share falls substantially for Network Health at the start of 2012, while rising sharply in all other plans. The patterns in both of these graphs reflect a combination of selection and within-enrollee cost reductions.

Figure 2. Adverse Selection Test: Past Patients at Partners Facilities

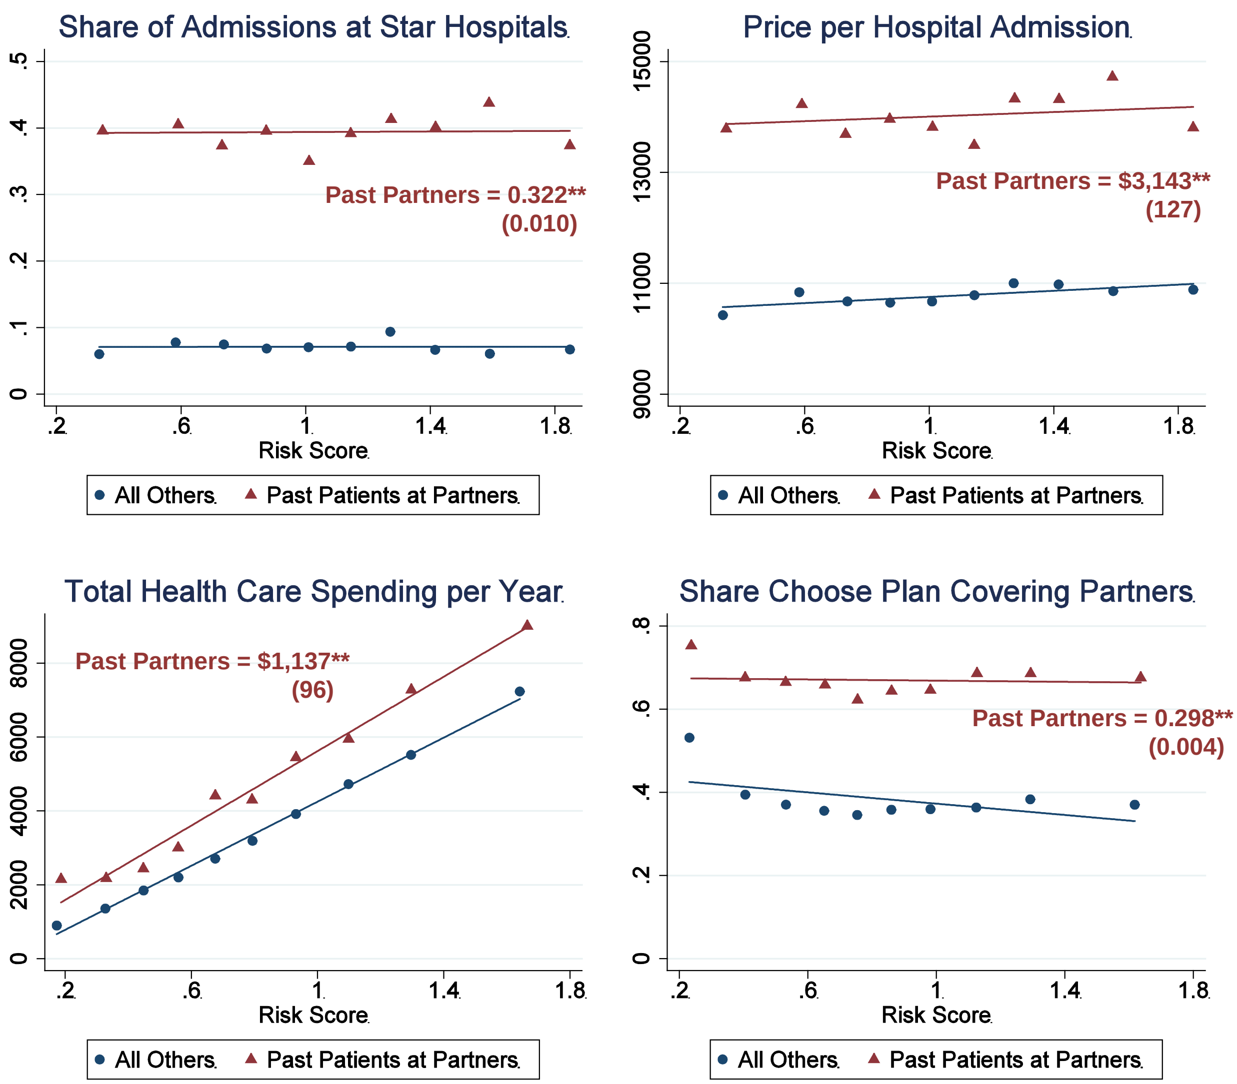

NOTE: The figures show binned scatter plots for results of the unused observables test for adverse selection, as described in Section 3.1. The figures compare outcomes for past Partners patients (in red triangles) to all other enrollees (blue circles), within bins of medical risk score (the x-axis). The lines are best-fit lines for each group. The data are at the individual x plan choice instance level for 2011-2013, the period during which I have full risk adjustment data. Cost and hospital use outcomes are defined as averages over the subsequent year. "Past Patients at Partners" is a dummy for whether an individual is observed using a Partners facility for outpatient care prior to the given plan choice timing. The sample excludes new enrollees in the exchange (for whom past utilization data is not observed). The bottom right figure further limits the sample to "re-enrollees" (who rejoin the exchange after a gap in coverage) who I can be sure are making active choices. Each figure reports the key "unused observable" coefficient on past Partners patient status and its standard error (clustered at the individual level). The full regressions are reported in Appendix Table 2.

Figure 3. Plan Switching and Selection when Network Health Drops Partners in 2012

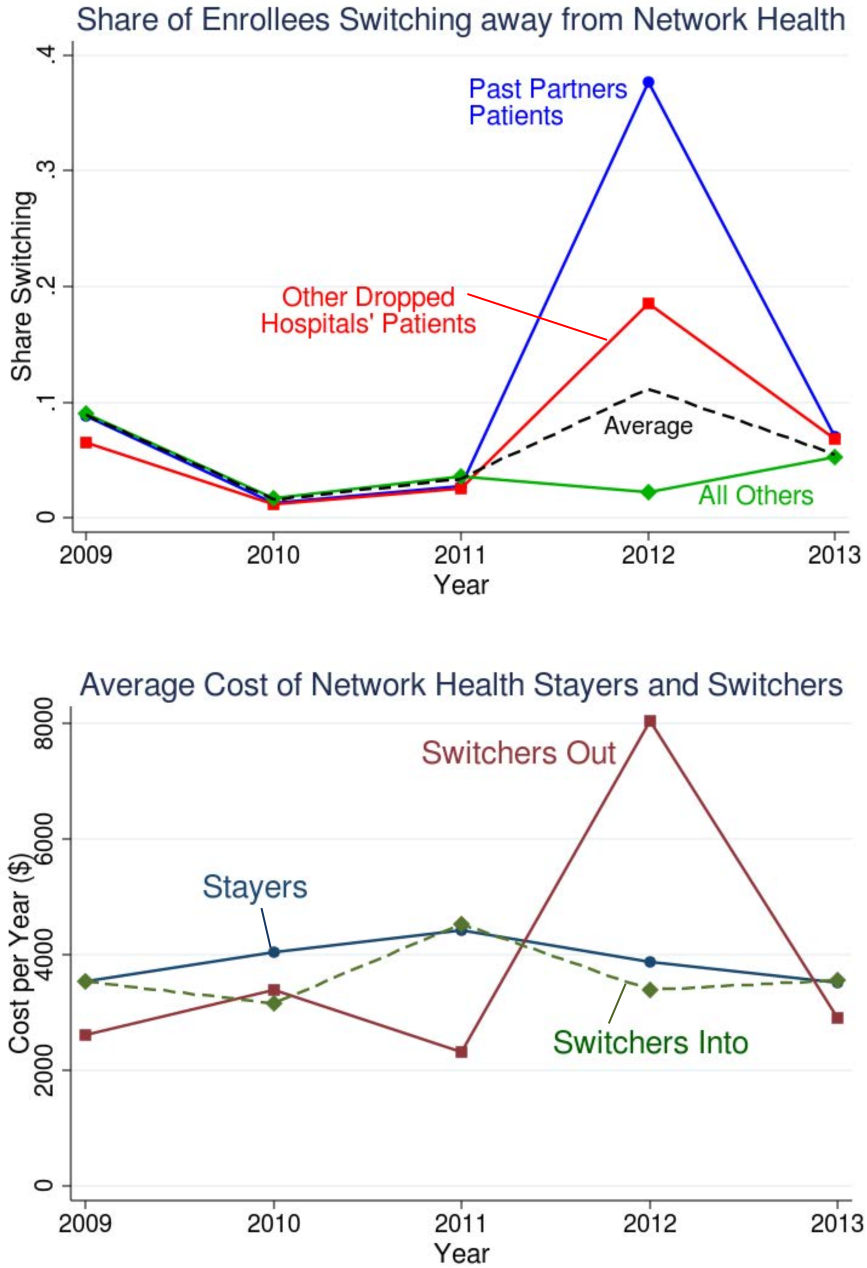

NOTE: These figures show switching and selection patterns for Network Health enrollees around its 2012 dropping of Partners and several other hospitals. The top graph shows the share of Network Health's enrollees who switch out of the plan at each year's open enrollment – separately for past Partners patients (in blue), past patients of other dropped hospitals (in red), all other enrollees (in green), and the average for all enrollees (black dashed line). Past patients of Partners and other dropped hospitals are much more likely to switch away from the plan in 2012 than other years. The bottom graph shows the average cost (during the prior year) of stayers, switchers out of, and switchers into Network Health at each year's open enrollment. In 2012, the cost of switchers out of Network Health was sharply higher, driven by the exit of past Partners patients. These results also hold for risk-adjusted costs (see Appendix Table 4 for additional details on the cost of stayers and switchers).

Figure 4. Cost Reductions for Stayers in Network Health

Panel A: Overall Effects

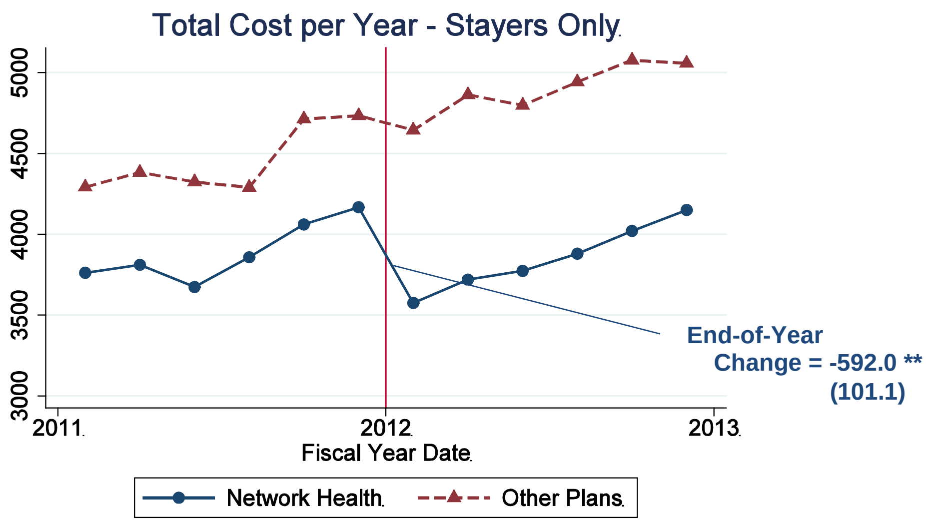

Panel B: Heterogeneity in Effects

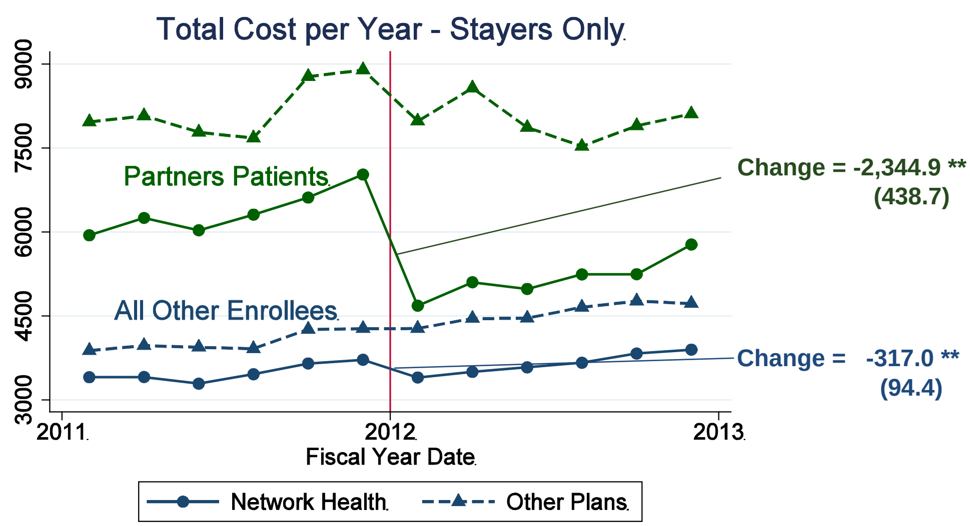

NOTE: These graphs show annualized costs for enrollees who stayed in Network Health around its dropping of Partners at the start of 2012. The estimates are based on panel regressions with individual fixed effects, to account for the unbalanced panel (due to churn in and out of the market). See Section 3.3 for the regression specification. The points plotted are regression coefficients, with the mean for each group in the last period of 2011 added to all points for the series. Panel A shows the effects for stayers in Network Health (blue solid lines), compared to stayers in other plans (red dashed lines), which did not make major network changes at this time. Panel B further breaks this down into past Partner patients (in green) and all other enrollees (in blue), with solid lines continuing to denote Network Health and dashed lines other plans.

Figure 5. Premium Coefficient Identification: Test of Parallel Trends Assumption

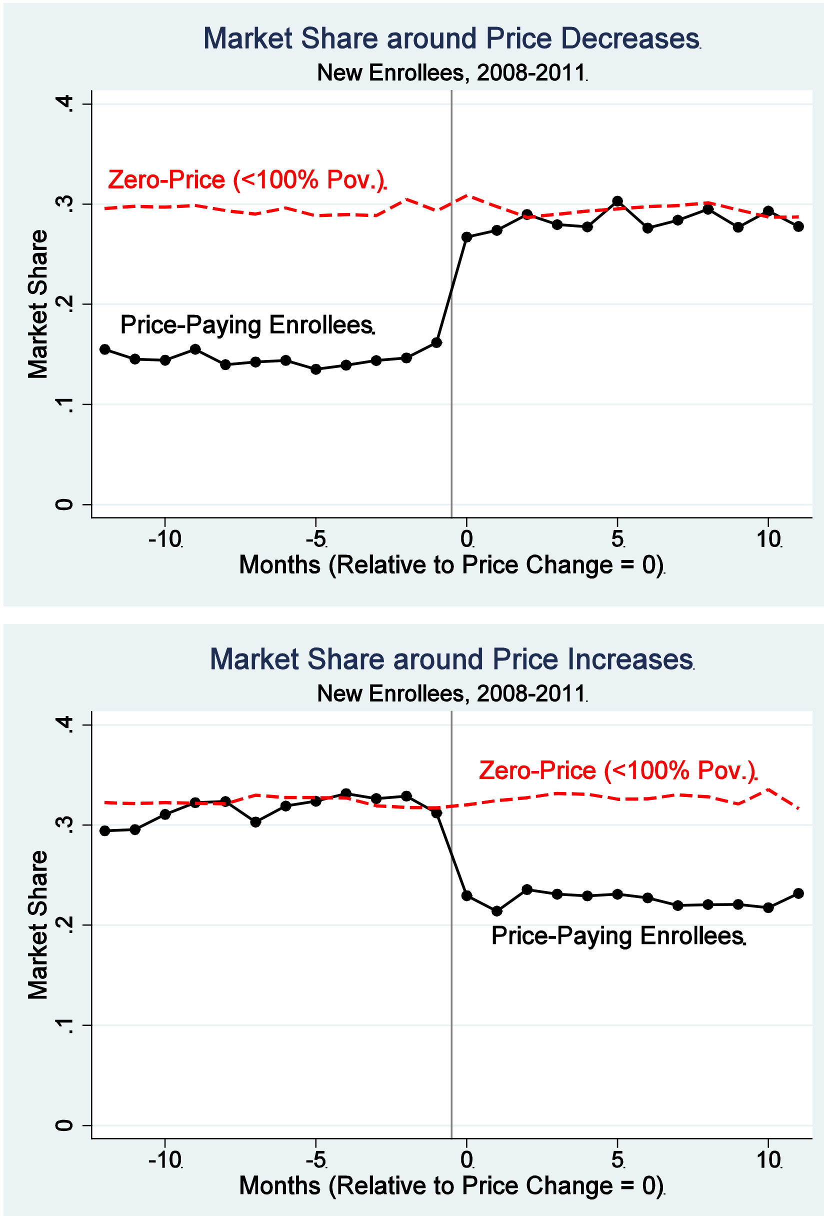

NOTE: These graphs show the source of identification for the premium coefficients in plan demand and test the key parallel trends assumption for the difference-in-differences approach. Each graph shows average monthly plan market shares among new enrollees for plans that at time 0 decreased their prices (top figure) or increased their prices (bottom figure). Each point represents an average market share for an independent set of new enrollees. The identification comes from comparing demand changes for above-poverty price-paying enrollees (for whom premium changes at time 0) versus below-poverty zero-price enrollees (for whom premiums are unchanged at $0). Consistent with the parallel trends assumption, trends in shares are flat and parallel for both groups at times other than the premium change but change sharply for price-payers only at the price change. The sample is limited to fiscal years 2008-2011. I drop 2012+ because below-poverty new enrollees became subject to a limited choice policy that required them to choose lower-price plans. In the demand model, I include dummy variables for below-poverty enrollees in 2012+ so that I do not use their market shares for identification.

Table 1. Hospital Price Estimates: Most Expensive Hospitals

	Hospital	System	Average Payment (per patient)	Hospital Price Model: Price (severity-adj.)	Hospital Price Model: Avg. Severity
			(1)	(2)	(3)
1	Brigham & Women's	Partners	$23,278	$20,474	1.12
2	Mass. General (MGH)	Partners	$21,428	$19,550	1.09
3	Boston Med. Ctr. (BMC)	---	$16,850	$15,919	1.05
4	Charlton Memorial	Southcoast	$14,411	$14,210	1.03
5	Umass Med. Ctr.	UMass	$14,941	$14,111	1.07
6	Tufts Med. Ctr.	Tufts/NEQCA	$15,328	$14,038	1.10
7	Baystate Med. Ctr.	Baystate	$13,715	$12,223	1.11
8	St. Luke's	Southcoast	$11,786	$12,113	0.97
9	Beth Israel Deaconess	CareGroup	$12,971	$11,787	1.08
10	Tobey Hospital	Southcoast	$11,427	$11,777	0.97
	All Other Hospitals	---	*$8,267*	*$8,549*	*0.96*

NOTE: This table shows the most expensive hospitals in my CommCare data, ranked by the hospital price measure in column (2). All measures are averages over in-network hospital admissions over fiscal years 2008-2013. Column (1) shows the average insurer payment (winsorized at $150,000 per admission to remove extreme outliers). Columns (2)-(3) shows average severity-adjusted prices and patient severity, estimated from a model described in Section 4.1.

Table 2. Coverage of Partners Hospitals by CommCare Plans

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

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

Table 3. Hospital Choices Model Estimates

VARIABLE	Coeff.	Std. Error
Distance to Hospital:		
Distance in Miles (avg. coeff.)	-0.160***	(0.001)
Distance^2 (avg. coeff.)	0.001***	(1e-5)
Distance Interactions:		
x Income > Poverty	-0.0064***	(0.0006)
x Age / 10	-0.0035***	(0.0002)
x Severity Weight	-0.0030*	(0.0010)
x Emergency	-0.0141***	(0.0006)
Out-of-Network Disutility		
Out-of-Network x BMC	-1.265***	(0.033)
Out-of-Network x CeltiCare	-1.601***	(0.056)
Out-of-Network x Fallon	-1.338***	(0.057)
Out-of-Network x NHP	-0.426***	(0.047)
Out-of-Network x Network	-1.133***	(0.035)
Out-of-Network x Emergency	0.010	(0.033)
Past Use of this Hospital		
Outpatient Care	2.182***	(0.013)
Hospital Characteristics		
Hospital Dummies	Yes	
Severity x Academic Med. Ctr. (avg).	2.099***	(0.043)
Severity x Teaching Hosp	1.017***	(0.044)
Diagnoses x Hospital Services (largest coeffs.):		
Mental: Psych. Services	1.837***	(0.038)
Pregnancy: Obstetrics Services	1.122***	(0.073)
Injury: Level 1 Trauma Center	0.793***	(0.036)
Cancer: Oncology Services	0.710***	(0.082)
Model Statistics:		
Pseudo-R^2 (McFadden's)	0.534	
R^2 in Shares (Area-Plan-Yr Level)	0.703	
Num. Choice Instances	74,308	

Std. Errors in parentheses. * = 5% sign., ** = 1% sign., *** = 0.1% sign.

NOTE: The table shows estimates for the multinomial logit hospital choice model described in Section 4.1. The coefficients shown are interpretable as entering the utility function describing hospital choice. Past use variables are dummies for whether a patient has previously used each specific hospital (before the current plan year and at least 30 days before the current admission). Severity is an estimated summary measure from the hospital price model described in Appendix B. In addition to the variables shown, the model includes: distance interacted with exchange region, detailed income group (by 50% of poverty), and gender; severity interacted with separate dummies for each academic medical center; and five additional diagnosis x hospital service interactions (circulatory diagnosis interacted with cath lab, interventional cardiology, and heart surgery services; pregnancy diagnosis x NICU; and musculoskeletal diagnosis x arthritis services).

Table 4. Insurance Plan Demand Estimates

VARIABLE	Coeff.	Std. Error	
Premium: Avg. Coeff. (*normalized*)	**-1.000*** **	**(0.025)**	
x 0-100% Poverty -- *Omitted (no prems.)*	---		
x 100-150% Poverty	-1.239***	(0.038)	
x 150-200% Poverty	-0.829***	(0.024)	
x 200-250% Poverty	-0.610***	(0.015)	
x 250-300% Poverty	-0.557***	(0.016)	
x Age/5 (average effect)	0.035***	(0.002)	
Hospital Network			
Network Utility x <100% Poverty	6.227***	(0.963)	
Network Utility x 100-150% Poverty	7.352***	(0.997)	
Network Utility x 150-200% Poverty	7.554***	(1.006)	
Network Utility x 200-250% Poverty	7.856***	(1.344)	
Network Utility x 250-300% Poverty	8.373***	(1.930)	
Past-Used Hospitals Covered (share)	5.736***	(0.853)	
x Past-Used Partners Hospitals	11.546***	(0.771)	
Switching and Inertia			Passive Prob.
Average Inertia Coefficient	95.638***	(0.234)	94.6%
x Drops Past-Used Hospital (Non-Partners)	-27.275***	(1.010)	71.1%
x Drops Past-Used Hospital (Partners)	-47.493***	(0.960)	51.8%
Plan Brand Effects (average)			
BMC HealthNet (normalized)	0.000	---	
CeltiCare	-22.644***	(0.953)	
Fallon	11.285***	(0.875)	
Neighborhood Health Plan	-2.742***	(0.300)	
Network Health	-3.746***	(0.334)	
Model Statistics			
R^2 in Share (Area-Income-Year)	0.962		
Model w/ Only Avg. Plan Dummies	0.872		
No. Choice Instances	1,588,889		
No. Unique Individuals	611,455		

* = 5% sign., ** = 1% sign., *** = 0.1% sign.

NOTE: This table shows estimates for the multinomial logit plan choice model described in Section 4.1. Premium is the monthly price paid by consumers, which typically varies by $20-60 across plans. (In addition to the interactions shown, the full model contains premium interactions with 5-year age groups and gender.) I normalize the average consumer's premium coefficient to -1.0, so all other coefficients are interpretable as dollar values. Network utility is the consumer-specific expected utility measure for a plan's hospital network, defined in Appendix B. Past-used hospitals coverage is the share of an enrollee's previously used hospitals that a plan covers, with a separate interaction for the star Partners hospitals. Switching and inertia are coefficients on a dummy variable for the current plan. The coefficients are interpretable as "switching costs" in dollars per month; the passive probabilities are the implied share of enrollees who passively stick with their current plan (as derived in Appendix B). The plan brand effects are coefficients on dummies for each plan. I show average values; the full model contains region-year- and region-income group-specific plan dummies.

Table 5. Analysis of Selection: Relationship between Value and Cost of Partners Coverage

Cons. Valuation of Partners Coverage		**Insurer Costs ($/month)**				
		Costs without Partners Coverage		**Cost Increase from Partners Covg.**		
Percentile of WTP	WTP ($/month)	Unadjusted	Risk Adj.	ΔCost	%Δ	ΔCost - Partners Markup
	(1)	(2)	(3)	(4)	(5)	(6)
< 50%	$0.5	$297.0	$304.2	$8.4	2.8%	$7.2
50-70%	$2.6	$278.0	$300.5	$16.0	5.7%	$11.5
70-80%	$5.0	$289.5	$304.2	$20.6	7.1%	$14.0
80-90%	$9.1	$327.9	$319.0	$28.5	8.7%	$16.6
90-95%	$19.8	$417.0	$348.8	$45.3	10.9%	$29.9
95-98%	$32.5	$455.3	$340.1	$55.9	12.3%	$38.0
98-100%	$52.8	$468.6	$356.6	$63.9	13.6%	$39.7
Average	***$5.9***	***$311.9***	***$310.0***	***$18.4***	***5.9%***	***$12.9***

NOTE: This table shows the model's implication for the relationship between enrollees' costs and their value coverage of the star Partners hospitals – the key relationship driving adverse selection. As described in the text, the sample is all enrollees in 2012 and costs and values are calculated for a single plan (Network Health). Figures for other years and plans are qualitatively similar. The table shows averages for consumer groups, sorted by increasing percentiles of willingness-to-pay (WTP) for Partners coverage. Column (1) shows each group's WTP (per month) of Partners coverage. Columns (2)-(3) show insurer costs for the group if the plan does *not* cover Partners (both unadjusted and after the exchange's risk adjustment). These are increasing in WTP, indicating selection on cost levels (both before and after risk adjustment). Column (4)-(5) report statistics on each group's cost increase when the insurer covers Partners. These are also increasing in WTP, indicating selection on cost increases (moral hazard). The final column (6) subtracts a measure of Partners' hospital markup to approximate the true social cost increase from covering Partners.

Table 6. Equilibrium Simulation Results

Panel A: Mass. Exchange Population and Policies

Source	Year	Variable	Insurance Plan: BMC	Fallon	Network Health	NHP	CeltiCare
Observed	2011	Partners Covg.	No	No	Yes	Yes	Yes
		Price*	$424.6	$425.7	$422.6	$425.7	$404.9
Simulated	2011	Partners Covg.	No	No	No	No	Yes
		Price*	$425.7	$425.7	$425.7	$425.7	$404.9

** Exchange imposed maximum price of $425.7 and minimum price of $404.9*

Panel B: ACA-Like Population and Policies

Source	Year	Variable	Insurance Plan: BMC	Fallon	Network Health	NHP	CeltiCare
Simulated	2011	Partners Covg.	No	No	No	No	No
		Price	$411.9	$397.4	$389.7	$413.0	$322.1
Simulated	2012	Partners Covg.	No	No	No	No	No
		Price	$422.1	$437.6	$367.7	$417.9	$358.2
Simulated	2013	Partners Covg.	No	No	No	No	No
		Price	$421.1	$477.4	$418.2	$447.3	$400.0

Panel C: ACA-Like, Partners Hospital Markups Eliminated

Source	Year	Variable	Insurance Plan: BMC	Fallon	Network Health	NHP	CeltiCare
Simulated	2011	Partners Covg.	No	No	No	No	No
		Price	$411.9	$397.4	$389.7	$413.0	$322.1
Simulated	2012	Partners Covg.	No	No	No	No	No
		Price	$422.1	$437.6	$367.7	$417.9	$358.2
Simulated	2013	Partners Covg.	No	No	No	Yes	No
		Price	$420.6	$478.0	$417.8	$466.3	$400.0

NOTE: These tables show equilibrium results for the insurance market simulations described in Section 5.1. The tables show their equilibrium choices of Partners coverage and price. Panel A shows simulations using the Massachusetts exchange's actual enrollee population and policies for 2011 and compares simulated coverage and prices to the observed values. (I do this comparison only for 2011 because of complications with simulating equilibrium for other years.) Panel B conducts simulations with a population and policies similar to those in the ACA exchanges, as described in the text. In both Panels A and B, Partners hospital prices are held fixed at levels estimated from the CommCare data. In Panel C, I simulate the market with the ACA-like population and policies, but with Partners' hospital prices reduced to a measure of their average costs (i.e., their markup over average cost is eliminated).

Table 7. Counterfactual Policy Simulations

Modified Risk Adjustment

Over-Adjustment Factor	Plan Statistics			Welfare Analysis (per member-month)				
	Covering Partners	Minimum Price	Avg. Price Other Plans	ΔCons. Surplus	Insurer Profit	Partners Net Rev.	Govt. Costs	ΔSocial Surplus
None	None	\$358.2	\$411.3	\$0.0	\$25.7	\$0.1	\$318.5	\$0.0
25%	None	\$369.7	\$415.4	\$3.4	\$29.4	\$0.2	\$326.1	-\$0.5
50%	None	\$377.3	\$417.0	\$7.1	\$32.5	\$0.2	\$333.2	-\$0.8
75%	None	\$382.9	\$418.5	\$9.9	\$34.6	\$0.2	\$338.2	-\$0.9
100% (eq 1)	NHP	\$386.3	\$424.9	\$9.0	\$37.3	\$1.6	\$341.9	-\$1.3
(eq 2)	CeltiCare, NHP, Network	\$409.3	\$438.9	\$17.9	\$42.5	\$4.5	\$363.3	-\$5.8

Subsidy for Partners Coverage

Subsidy (per member-month)	Plan Statistics			Welfare Analysis (per member-month)				
	Covering Partners	Minimum Price	Avg. Price Other Plans	ΔCons. Surplus	Insurer Profit	Partners Net Rev.	Govt. Costs	ΔSocial Surplus
None	None	\$358.2	\$411.3	\$0.0	\$25.7	\$0.1	\$318.5	\$0.0
\$4	BMC	\$358.1	\$412.6	-\$0.4	\$25.7	\$0.7	\$319.4	-\$0.7
\$8	BMC	\$358.0	\$411.8	-\$0.1	\$26.1	\$0.8	\$320.2	-\$0.7
\$12	BMC	\$358.0	\$411.0	\$0.4	\$26.5	\$0.8	\$321.1	-\$0.8
\$16	BMC, CeltiCare, NHP	\$368.5	\$411.9	\$10.1	\$28.9	\$3.2	\$337.9	-\$3.0
\$20	All Plans	\$363.2	\$406.1	\$12.8	\$26.6	\$5.2	\$342.9	-\$5.7

NOTE: These tables show results of counterfactual policy simulations, as discussed in Section 5.3. The top table shows simulations that modify risk adjustment to "over adjust" for observable risk. The bottom table shows simulations where the exchange provides an extra subsidy to any plan that covers Partners. All simulations are for the ACA-like population and policies in 2012, so the baseline results (top row of each table) are identical to the 2012 equilibrium in Table 6. Each table lists which plans cover Partners, the minimum plan price, and average price of all other plans. They also list welfare statistics (in dollars per member-month): the change in consumer surplus (with the baseline normalized to \$0), insurer profit, Partners' net inpatient hospital revenue, and government subsidy costs. The final column show the change in social surplus (relative to the baseline), which is defined as the sum of consumer surplus, insurer profit, and Partners net revenue, minus government costs.

Appendix A. Basic Theory

In this appendix, I present a simple model to illustrate how coverage of expensive star hospitals can lead to adverse selection, even with sophisticated risk adjustment in place. Adverse selection occurs when consumers with high value for generous insurance also tend to have high unobserved (or unpriced) costs. The literature has typically equated higher costs with greater *medical risk* – i.e., that higher-cost consumers are sicker. Key to my model is a second, conceptually different source of cost heterogeneity: *preferences* for using expensive providers when sick. While the model focuses on expensive star hospitals, the theory applies more broadly to preferences for any high-cost treatment option (e.g., branded vs. generic drugs, or high- vs. low-cost procedures). Because the insurer covers all or part of the excess cost of the expensive option, people who are more likely to use it are higher cost to the insurer.

Simple Model

Consider a model where insurers compete on prices and a single generosity choice: whether to cover a star hospital, *S*, in its network. For simplicity, assume that the star hospital's price is a uniform τ_S per visit for all insurers.[58] All other "non-star" hospitals charge $\tau_{NS} < \tau_S$ per visit and are covered by all insurers. Importantly, insurers that cover *S* do not fully pass on its higher price to patients but instead cover the price differential. Here, for simplicity, I assume patient fees (copays) are zero.[59]

After seeing insurers' offerings, consumers choose a plan and when sick, choose among in-network hospitals. Consumers vary in two ways:

1. **Medical risk**, $r_{i,d}$, for various diagnoses $d = 1,\ldots,D$
2. **Value for the star hospital**, $v_{i,d}^S$, for each diagnosis *d*

Medical risk equals a consumer's probability of being hospitalized for diagnosis *d*, which I model as an exogenous event. Value for the star hospital (or what I call "preferences") is consumers' diagnosis-specific willingness-to-pay for the star hospital relative to the next best alternative. This value can be negative if a non-star hospital is preferred (e.g., because of greater convenience). Let $I_{i,d}^S$ indicate whether the consumer chooses the star hospital for diagnosis *d* if covered. Assume that consumers do not use the star hospital if out of network. Define individuals' overall risk as $r_i \equiv \sum_d r_{i,d}$, and the share of

[58] This and many other assumptions are made for presentational simplicity and are relaxed in the structural model.

[59] If they were non-zero, τ_S and τ_{NS} would equal the insurer's net cost (= hospital price – patient copay). The assumption that insurers cover part of the fee differential ensures that $\tau_S > \tau_{NS}$.

illnesses for which they choose the star hospital as $s_i \equiv \frac{1}{r_i}\sum_d r_{i,d} I^S_{i,d}$. Finally, $\Delta\tau \equiv \tau_S - \tau_{NS}$. Expected costs for consumer *i* in a plan that does not cover *S* equal:

$$C_i^{NoCover} = r_i \cdot \tau_{NS} \tag{8}$$

while costs in a plan that covers *S* equal:

$$\begin{aligned} C_i^{CoverS} &= \quad r_i \cdot \tau_{NS} \quad + \; r_i \cdot s_i \cdot \Delta\tau \\ &\equiv \quad C_i^{NoCover} \quad + \quad \Delta C_i \end{aligned} \tag{9}$$

This formula shows the two sources of cost variation: illness risk (r_i) and likelihood to choose the star hospital when sick (s_i). Although these may be correlated – sicker people may be more likely to choose star hospitals – these are conceptually separate drivers of costs. A key distinction is that high-s_i types are more expensive only in plans that cover the star hospital they prefer. Preference for the star hospital therefore affects enrollees' *cost differences* (ΔC_i) across plans – often called the moral hazard effect of covering *S*.[60] This heterogeneity in cost differences has implications for the nature of selection and the effectiveness of risk adjustment, as I discuss below.

Prior to realizing health shocks, consumers choose among plans based on plans' prices and coverage of hospital *S*. Let the utility of a plan not covering *S* be normalized to zero. Assume that consumers' extra utility for a plan that covers *S* equals their *ex-ante* expected value of access to *S*, or:

$$U_i^{CoverS} = \sum_d r_{i,d} I^S_{i,d} v^S_{i,d} = r_i \cdot s_i \cdot \overline{v}_i^S \tag{10}$$

where $\overline{v}_i^S \equiv \frac{1}{r_i s_i}\sum_d r_{i,d} I^S_{i,d}$ is the consumer's average value for the star hospital conditional on use. The key feature of this assumption is that consumers' utility for a plan covering the star hospital is *linked to their likelihood of using it* $(= r_i \cdot s_i)$. This link – which is built into the standard "option demand" model of Capps et al. (2003) – generates the correlation between demand and costs that drives adverse selection.

Assume that each plan *j* sets a single price P_j that cannot vary across consumers.[61] Although prices cannot vary, the exchange risk adjusts payments based on consumer observables Z_i so a plan in total receives $P_j + RA(Z_i)$ for consumer *i*.[62] The risk adjustment function is set to offset a consumer's expected extra costs, so $RA(Z_i) = E(\mathrm{C}_{ij} \mid Z_i) - \overline{C}$ (where $\overline{C}$ is overall average cost). If risk adjustment captured

[60] In the health insurance literature, "moral hazard" typically refers to changes in enrollee's utilization in response to more generous insurance. Even though not "hidden action" in the contract theory sense, the term is applied because the change in action is not contracted on, often because of regulatory constraints.

[61] Assume that any subsidies are a flat amount so that consumer premium differences are equal to price differences.

[62] For expositional simplicity, this formulation follows the ACA's risk adjustment method. In the structural model, I follow Massachusetts' multiplicative risk adjustment method, with a risk score multiplying the plan's price.

costs perfectly, a plan's profit margin would be a constant $P_j - \overline{C}$ for all consumers. However, risk adjustment is unlikely to offset the higher costs of high-s_i types for two reasons. First, the standard risk adjusters in Z_i (typically age, sex, and medical diagnoses) are intended to capture medical risk, not hospital choices – though, in principle hospital choice predictors could be added. Second, and more fundamentally, a single risk adjustment value $RA(Z_i)$ cannot offset the heterogeneity in cross-plan cost differences (moral hazard) that occurs in this setting (a point emphasized by Einav et al. 2015). Costs vary not only because of consumer heterogeneity but because of the *interaction* of consumer types with the hospitals a plan covers.

Implications for Market Equilibrium

This model has several implications for market equilibrium, which I discuss in turn. For simplicity, I continue to assume a setting where there are two otherwise identical plans: one that covers the star hospital *S* and one that does not. Define $\Delta P = P_{CoverS} - P_{NoCover}$ as the price difference between these plans. Note that the set of enrollees who choose the plan covering *S* are those for whom the utility exceeds this price differential: $U_i^{CoverS} \geq \Delta P$.

(a) Selection on two dimensions of costs: Adverse selection occurs if plans that cover the star hospital tend to attract enrollees with high risk-adjusted costs. This selection can occur through two cost dimensions: unobserved risk and the cost difference from covering *S*. To see this formally, assume that the risk adjustment formula, $RA(Z_i)$, is based on costs in plans not covering *S*. Define $e_i \equiv C_i^0 - RA(Z_i)$ as the error in this prediction. Define average risk-adjusted costs for enrollees in plan *j* as $AC_j = E\left(C_{ij} - RA(Z_i) \mid i \text{ chooses } j\right)$. With a bit of algebra, the difference in risk-adjusted average costs between the two plans, $\Delta AC \equiv AC_{CoverS} - AC_{NoCoverS}$, can be decomposed as follows:

$$\Delta AC = \underbrace{\overline{\Delta C}}_{\text{(1) Cost Difference for Avg. Person}} + \underbrace{\left(\overline{e}^{CoverS} - \overline{e}^{NoCoverS}\right)}_{\text{(2) Selection on Unobs. Risk}} + \underbrace{E\left(\Delta C_i - \overline{\Delta C} \mid U_i^{CoverS} \geq \Delta P\right)}_{\text{(3) Selection on Cost Difference}} \qquad (11)$$

where $\overline{e}^{CoverS}$ and $\overline{e}^{NoCoverS}$ are average risk adjustment errors for each plan type. Equation (11) separates out three components of average cost differences between plans. First, term (1) captures that a plan covering *S* has higher costs even for an average person because of the moral hazard effect of covering *S*. Term (2) captures traditional selection on unobserved risk. Without additional assumptions, the sign of this term is ambiguous. Whether people who like the star hospital are unobservably sicker or healthier is driven by context-specific factors that are not obvious *a priori*.

Finally, term (3) captures *selection on cost differences* (or selection on moral hazard). Unlike unobserved risk, there is a simple theoretical reason to expect a positive sign (adverse selection) for this term. The people who select plans covering *S* are those with $U_i^{CoverS} = r_i \cdot s_i \cdot \overline{v}_i^S \geq \Delta P$. Meanwhile, the cost difference is $\Delta C_i = r_i \cdot s_i \cdot \Delta\tau$. Because use of the star hospital ($= r_i \cdot s_i$) appears in both terms, the two are likely to be positively correlated.[63] Intuitively, propensity to use the star hospital drives both plan preferences and the cost difference between plans.

(b) Inefficient sorting across plans: To sort consumers efficiently, it is optimal for premium differences, ΔP, to equal *individual-specific* cost differences between plans, ΔC_i. In a model with homogenous cost differences ($\Delta C_i = \overline{\Delta C}$ for all *i*), this optimum would be attainable. The goal of risk adjustment in such a model is to eliminate selection on unobserved risk, so that in competitive equilibrium, $\Delta P = \Delta AC = \overline{\Delta C}$. This is the basic intuition underlying traditional risk adjustment.

With heterogeneity in ΔC_i, first-best sorting is unattainable with homogenous premium differences between plans – a point that has been emphasized by Bundorf et al. (2012). It is optimal to choose a plan covering *S* if and only if $U_i^{CoverS} \geq \Delta C_i$, which simplifies to $\overline{v}_i^S \geq \Delta\tau$.[64] But consumers choose it if $r_i s_i \overline{v}_i^S \geq \Delta P$. The discrepancy between these conditions leads to both errors of over- and under-purchase of plans covering the star hospital.

Even if the first-best is unattainable, it is interesting to ask how selection affects prices relative to a second-best optimal single premium difference. The second best is defined by the condition $\Delta P = E\left(\Delta C_i \mid U_i^{CoverS} = \Delta P\right)$, which equates price to the marginal enrollees' cost difference. Equation (11) shows that in a competitive equilibrium with $\Delta P = \Delta AC$, adverse selection on both unobserved risk and moral hazard pushes ΔP above this optimum. The intuition for unobserved risk is standard. For selection on moral hazard, the intuition is that the marginal type uses the star hospital less than the average person in the *S*-covering plan. The need to pool with these high-ΔC_i types discourages some people for whom access to *S* would be efficient.

(c) Star hospital coverage and market power: Adverse selection (through either channel) has a natural effect on insurers' incentives to cover the star hospital, and in turn on its market power in price negotiations. To study these issues, suppose that instead of perfect competition, there is an imperfectly

[63] As is evident, this positive correlation is not mathematically certain and could be broken if $\overline{v}_i^S$ were strongly negatively correlated with $r_i s_i$ -- i.e., if the people most likely to use the star hospital only slightly preferred it over other alternatives. While mathematically possible, this case seems unlikely.

[64] These conditions would be different if $\Delta\tau$ includes a hospital price markup, an issue I return to below.

competitive insurance market where each insurer bargains with the star hospital over its payment rate, τ_S, and inclusion in network. Assume that the star and non-star hospitals have marginal costs of mc_s and mc_{NS}, and that because of hospital competition $\tau_{NS} = mc_{NS}$. I do not specify a full bargaining model for the determination of τ_S but note that in standard models (e.g., Nash bargaining), a key determinant is an insurer's change in profits from shifting from not covering to covering *S* at a given τ_S, or:

$$\Delta\pi_j\left(\tau_S\right) = \left[\Delta P_j - \Delta AC_j\left(\tau_S\right)\right] \cdot Q_j^{CoverS} + \left[P_j^{NoCover} - AC_j^{NoCover}\right] \cdot \Delta Q_j \tag{12}$$

where all of these terms are equilibrium values, which incorporate the shift in plan prices when plan *j* adds *S* to its network.[65] Adverse selection implies a larger increase in average costs (ΔAC_j) when a plan covers *S*. This makes covering the star hospital less profitable at any given payment rate τ_S.

This lower insurer profitability in turn affects the payment rate the star hospital can extract. Intuitively, adverse selection improves the insurer's *threat point* (profits if it excludes *S*) in a bargaining game. Two possible outcomes can result. If the star hospital's high prices reflect high markups, adverse selection can discipline these markups and lead to lower τ_S without any plans dropping it from network. Alternatively, if the star hospital's high payment rates reflect high marginal costs, insurers may find it profitable to drop *S* even at $\tau_S = mc_S$, resulting in less equilibrium coverage of the star hospital.

Thus, adverse selection can have important implications for both equilibrium coverage and prices of star hospitals. For tractability in my structural model, I will only consider the coverage channel – I hold hospital prices fixed and simulate insurers' decision to cover/exclude the star hospital. However, readers should keep in mind the broader conceptual point that adverse selection in insurance markets can discipline star hospitals' market power. This point is an important caveat to the typical logic that popular hospitals for which consumers have high "willingness to pay" have the strongest market power. In markets where insurers compete (as opposed to most employer insurance settings), a hospital's market power is related to insurers' *profitability* of covering it. Profitability depends both on how much covering the hospital increases a plan's demand (roughly analogous to willingness to pay) but also on *which* consumers it attracts. If covering it attracts high-cost, unprofitable consumers, that hospital may have significantly less leverage to negotiate high prices.

[65] Depending on the timing of the game, this condition may implicitly include the equilibrium pricing response of other insurers' in the definition of quantities and average costs.

Appendix B. Structural Model and Estimation Details

B.1. Hospital Network Utility Definition

To generate a measure of network utility for plan demand, I follow the method of Capps et al. (2003). I define network utility based on the expected utility metric from the hospital demand system. Conditional on needing to be hospitalized, a consumer's utility of access to network $N_{j,t}$ in plan *j* is:

$$\begin{aligned} HospEU_{i,d,t,j}\left(N_{j,t}\right) &\equiv E\max_{h}\left\{\hat{u}_{i,d,t,j,h}\left(N_{j,t}\right)+\varepsilon_{i,d,t,j,h}\right\} \\ &= \log\left(\sum_{h}\exp\left(\hat{u}_{i,d,t,j,h}\left(N_{j,t}\right)\right)\right) \end{aligned} \tag{13}$$

where $\hat{u}_{i,d,t,j,h}\left(N_{j,t}\right) \equiv u_{i,d,t,j,h} - \varepsilon_{i,d,t,j,h}$. At the time of plan choice, however, consumers do not know their hospital needs. Instead, they have expectations of their hospital use frequency for each diagnosis *d* over the coming year, which I denote $freq_{i,d,t}$. Given this expectation, the *ex-ante* expected network utility is:

$$NetworkUtil_{i,j,t}\left(N_{j,t}\right) \equiv \sum_{d} freq_{i,d,t} \cdot HospEU_{i,d,t,j}\left(N_{j,t}\right) \tag{14}$$

This network utility in (14) is what I include in plan demand. To calculate it, I first use my data to estimate a Poisson regression of the annual number of hospitalizations for each diagnosis on individuals' age and demographics.[66] I use the predicted values from these regressions for $freq_{i,d,t}$. Next, I calculate the value of $HospEU_{i,d,t,j}\left(N_{j,t}\right)$ for each plan and diagnosis, using the individual's location and demographics at the time of plan choice.[67] Finally, I input these values into equation (14) to calculate network utility. Because network utility does not have natural units, I normalize it so that 1.0 is the average decrease in utility for Boston-region residents when Network Health dropped Partners in 2012.

B.2. Moments for Insurance Choice Model Estimation

I estimate the plan demand model parameters by matching moments that fall into two categories. First, for plan dummies, I match market shares for the appropriate region/year/income group *g*. These market shares uniquely identify plan mean utilities, which in my case are equivalent to the plan dummies.[68] The formula for these market share moments is:

---

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

[67] The two unknown covariates at the time of plan choice are severity and emergency status. For emergency status, I use diagnosis-specific emergency probabilities to average the values of *HospEU* for each possibility. For severity, I use a predicted value from a regression of severity on age-sex groups and emergency status.

[68] A difference in my setting from the standard BLP approach is that I treat the plan dummies as parameters, with associated standard errors, since both they and the characteristics coefficients are estimated from a dataset of the same size (the full market data). In previous applications including Berry, Levinsohn, and Pakes (2004), the micro data came from a sample, while the market shares came from aggregate data on the whole market.

$$G_{j,g}^{(1)}(\theta)=\tfrac{1}{N}\sum_{i,j,t}1\{i,t\in g\}\cdot\left[1\{y_{it}=j\}-Pr(y_{it}=j\mid\theta)\right]$$

where θ is the parameter vector, $1\{y_{it}=j\}$ is an indicator for whether individual *i* chose plan *j* at time *t*, and $Pr(y_{it}=j\mid\theta)$ is the predicted choice share from the logit model.

Second, for the coefficients for premium, network utility, and other observed characteristics (which are interacted with observed enrollee attributes), I match the average values for chosen plans in the data to those in the model. Specifically, the moments for characteristic $X^{(k)}$ (e.g., premium) interacted with enrollee attribute $Z^{(r)}$ (e.g., income) are:

$$G_{k,r}^{(2)}(\theta)=\tfrac{1}{N}\sum_{i,j,t}X_{ijt}^{(k)}Z_{it}^{(r)}\cdot\left[1\{y_{it}=j\}-Pr(y_{it}=j\mid\theta)\right]$$

Another way of interpreting these is as matching the *covariance* between plan characteristics and household attributes. In the case of observing the full market, these moments are equivalent to the micro BLP covariance moments. These moments are also equivalent to first-order conditions from the associated maximum likelihood problem.

Stacking all of the moments into a vector $G(\theta)$, the MSM estimator searches for the parameter θ that minimizes the weighted sum of squared moments, $G(\theta)'\cdot W\cdot G(\theta)$. Because the system is just-identified, I am able to match the moments exactly, making the solution invariant to the choice of *W*. I calculate standard errors using the standard GMM sandwich formula.

B.3. Inattention Interpretation of Plan Choice Inertia Coefficients

For current enrollees, I included in the logit demand model a dummy variable for their current plan, so their full demand utility was:

$$U_{ijt}^{Curr}=\hat{U}_{ijt}+\underbrace{\chi(Z_i)\cdot 1\{j=CurrPlan\}}_{\text{Switching Cost / Inertia}}+\varepsilon_{ijt}^{Plan} \tag{15}$$

where $\hat{U}_{ijt}$ is the plan utility for new enrollees (defined in Section 4.1), excluding the ε_{ijt}^{Plan}. In this equation, $\chi(Z_i)$ is interpreted as a switching cost – an extra utility for the current plan needed to rationalize the low level of plan switching. The plan demand estimates in Table 4 reports these switching costs but also an alternate interpretation based on an inattention model. I show here how I derive the inattention/passive probability reported in Table 4.

Consider a two-step model in which the first step models whether enrollees make an active choice, and the second step models plan choice conditional on being active. The second step is standard and follows the logit model for new enrollees (or current enrollees excluding switching cost):

$$Pr\left(y_{it}=j\mid Active\right)=\frac{\exp\left(\hat{U}_{ijt}\right)}{\sum_{k}\exp\left(\hat{U}_{ikt}\right)}$$

The first step is a reduced form model of being passive:

$$Pr_{it}\left(Passive\right)=\frac{\exp\left(\hat{U}_{i,j_{curr},t}+\tilde{\chi}_i\right)}{\exp\left(\hat{U}_{i,j_{curr},t}+\tilde{\chi}_i\right)+\exp\left(I_{i,Active,t}\right)}\quad\text{where } I_{i,Active,t}=\log\left(\sum_{k}\exp\left(\hat{U}_{ikt}\right)\right)$$

Notice that it is the choice probability from a two-choice logit model, where the utility of being passive is the current plan utility plus a reduced-form inertia coefficient $\tilde{\chi}_i$ (which is different from the switching cost χ). The utility of being active is $I_{i,Active,t}$, which is the inclusive value (or expected utility) from the second-stage active choice model.

I claim that if $\tilde{\chi}_i=\log\left(\exp\left(\chi(Z_i)\right)-1\right)$, the switching cost and inattention models have identical predictions for choice probabilities. For the current plan, the inattention model predicts a probability that it is chosen of $Pr_{it}\left(Passive\right)+\left(1-Pr_{it}\left(Passive\right)\right)\cdot Pr\left(y_{it}=j_{curr}\mid Active\right)$, which simplifies to:

$$Pr\left(y_{it}=j_{curr}\right)=\frac{\exp\left(\hat{U}_{i,j_{curr},t}+\chi\left(Z_i\right)\right)}{\exp\left(\hat{U}_{i,j_{curr},t}+\chi\left(Z_i\right)\right)+\sum_{k\neq j_{curr}}\exp\left(\hat{U}_{ikt}\right)}$$

This equals the current plan's choice probability in the switching cost model in (15). Further, the inattention model's probability of switching to another plan *j* is $\left(1-Pr_{it}\left(Passive\right)\right)\cdot Pr\left(y_{it}=j\mid Active\right)$, which simplifies to:

$$Pr\left(y_{it}=j\right)=\frac{\exp\left(\hat{U}_{i,j,t}\right)}{\exp\left(\hat{U}_{i,j_{curr},t}+\chi\left(Z_i\right)\right)+\sum_{k\neq j_{curr}}\exp\left(\hat{U}_{ikt}\right)}$$

which is again equivalent to the choice probability from the switching cost model in (15).

Hence, these two models have equivalent predictions for choice probabilities. The plan demand results in Table 4 report both the average switching costs $\chi\left(Z_i\right)$ and the passive probability $Pr_{it}\left(Passive\right)$, as defined by the equation above.

B.4. Hospital Price Model

As discussed in Section 4.1, I estimate a model of hospital prices for a severity-adjusted admission. Because actual payment rules are unknown (and likely quite complicated), there is a need for simplification. I follow past work (e.g., Gowrisankaran et al. 2015) in estimating *average* payment factors that capture proportional differences across hospital-insurer pairs.[69] I estimate a Poisson regression (also known as a generalized linear model with a log link) of the form:

$$E\left[Payment_{i,j,h,t,a} \mid Diag_{ita}, Z_{ita}\right] = \exp\left(\rho_{j,h,t} + Diag_{ita}\lambda + Z_{ita}\gamma\right)$$

where a indexes the admission, $Diag_{ita}$ is the principal diagnosis, and Z_{ita} is other patient covariates. The key term is $\rho_{j,h,t}$, which is a coefficient that captures average payment differences across hospitals, insurers, and years. This effect is assumed to be proportional across all types of admissions, which is surely not exactly right but should capture a valid average effect.

For the principal diagnosis ($Diag_{ita}$), I use the Clinical Classification Software (CCS) dummies defined by the U.S. government's Agency for Healthcare Research and Quality. The additional covariates (Z_{ita}) include age, gender, income, and Elixhauser comorbidity dummies for the secondary diagnoses.

I specify a restricted model for $\rho_{j,h,t}$ to avoid over-fitting for hospital-insurer-year cells with small samples. Specifically, I start from the model:

$$\rho_{j,h,t} = \rho_{j,h,NetwStat(h,t)} + \rho_{j,Sys(h),t} + \rho_{j,t,NetwStat(h,t)}$$

The first term, $\rho_{j,h,NetwStat(h,t)}$, is a coefficient for each hospital-insurer pair, separately for in- versus out-of-network status. This term is constant across years except to the extent the hospital's network coverage changes. I include this term for all cells with at least 50 observations; otherwise, I omit it to avoid over-fitting (and capture prices for these smaller hospital-insurer pairs using the remaining terms). The second term, $\rho_{j,Sys(h),t}$, is a coefficient for each insurer-system-year triple for the top six hospital systems when covered in-network. This allows for a separate hospital price paths over time for each of the largest systems (including Partners). I do not include this term for hospitals in smaller systems or when a large system is out-of-network, with the exception that I always include these dummies for Partners regardless of whether it is in-network. The final term, $\rho_{j,t,NetwStat(h,t)}$, is a residual that allows for a separate effect for each plan, year, and network status. This captures the average insurer-specific price path for all smaller hospitals not included in one of the six largest systems.

I use the estimates of the equation above to define hospital prices as $\hat{P}_{j,h,t} \equiv \exp\left(\hat{\rho}_{j,h,t}\right)$ and an admission-specific severity measure as $\hat{\omega}_{i,t,a} \equiv \exp\left(Diag_{ita}\hat{\lambda} + Z_{ita}\hat{\gamma}\right)$. I scale $\hat{\omega}_{i,t,a}$ so that its mean is 1.0

[69] Following convention, I refer to these payment factors as "prices," although they are distinct from the actual negotiated prices. These payment factors capture both price differences and service quantity differences across hospitals (conditional on diagnosis) since both affect insurers' payment differences across hospitals.

and divide $\hat{P}_{j,h,t}$ by the same factor, so it can be interpreted as the hospital price for a patient of average severity. The average prices and severities for the 10 most expensive hospitals are shown in Table 1.

B.5. Non-Hospital Cost Model

For non-hospital costs (all costs except inpatient hospital care), I take a reduced form approach. I calculate monthly non-inpatient costs for each enrollee-year and use them to estimate the following Poisson regression model:

$$E\left(NonHospCost_{i,j,t}^{Obs} \mid Z_{it}\right) = \exp\left(\eta_{j(i),Reg(i),t} + Z_{it}\mu\right) \tag{16}$$

where Z_{it} are detailed enrollee diagnoses and demographics.[70] I use these estimates to define a region-year-specific plan effect $\nu_{j,Reg,t} \equiv \exp\left(\hat{\eta}_{j,Reg,t}\right)$. For counterfactual plan *k*, I assume that an enrollee's costs equal observed costs times the ratio of plan effects, $\left(\nu_{k,Reg,t} / \nu_{j,Reg,t}\right)$. This approach is clearly an approximation. However, the $\nu_{j,Reg,t}$ estimates should capture a valid average plan effect on costs absent unobserved cost-based selection into plans. Given that I have documented unobserved selection based on the exchange's risk adjustment, this assumption is clearly imperfect.[71] If there is residual selection, I will understate costs for plans attracting residually healthier enrollees and overstate costs in the opposite case. This will affect my estimates of the *level* of non-inpatient costs at observed networks but not the cost *difference* from network changes, which is the key statistic for welfare and plans' network choice.

Non-hospital costs are likely to change in some ways when Partners is added to network (e.g., if patients substitute to higher-price Partners doctors) but be unchanged in other ways (e.g., drug utilization is unlikely to change). To approximate the average effect of a plan adding/dropping Partners on non-inpatient costs, I use the following method. I first use the hospital cost model to calculate how inpatient costs change when Partners is added/dropped. I then assume that a plan's non-inpatient costs change in proportion to the *average* percent change in inpatient costs. I compute this average change separately by year, region, and past use of Partners, to allow for the effect of adding Partners to vary across enrollees.

Putting these steps together, the model for non-hospital costs for enrollee *i* in plan *j* at time *t* is:

$$c_{i,k,t}^{NonHosp}\left(N_{kt}\right) = \underbrace{NonHospCost_{i,j(i),t}^{Obs}}_{\text{Observed Costs}} \cdot \underbrace{\left(\frac{\nu_{k,Reg,t}}{\nu_{j(i),Reg,t}}\right)}_{\text{Ratio of Plan Effects}} \underbrace{\left(1 + \lambda \cdot \overline{\%\Delta HC}_{k,Reg,PastUse,t} \cdot 1\left\{k \text{ adds/drops Partners}\right\}\right)}_{\text{Partners Covg. Change Adjustment}}$$

---

[70] For diagnoses, I use the Hierarchical Condition Categories (HCC) defined by Medicare for its risk adjustment. I use HCCs observed in the *current* plan year so I can include diagnoses for new enrollees.

[71] The covariates in (16) will do somewhat better than the exchange risk adjustment because they include concurrently observed diagnoses, which allows for including diagnoses for new enrollees.

where the first terms are discussed above, and the final term is the adjustment if Partners is added/dropped. In this term, $\overline{\%\Delta HC}$ is the average percent change in hospital costs, and λ is a scaling factor capturing the average relationship between non-inpatient and inpatient costs. Based on a risk-adjusted regression at the plan-region-year-level, I estimate $\lambda = 0.38$.

Appendix C. Equilibrium Simulation Method Details

This appendix details the simple approach I use to incorporate a future profit effect in a static pricing model for my simulations in Section 5. Note that in a dynamic model, an insurer's pricing FOC includes a term capturing the effect of changing today's price on future profits earned from consumer *i*. I model this "future profit effect" as the product of the change in future demand $(\partial D_{ij}^{Fut} / \partial P_j)$ times an expected profit margin M_{ij}^{Fut}, which I assume is unaffected by today's price.

For the change in future demand, a lower price today increases today's demand and therefore increases the number of inertial enrollees in the future. To model this, I assume an exogenous, constant inertia probability ρ at each year's switching choice, which I set at 89%. I use 89% rather than the 95% inertia probability reported in the plan demand estimates based on a rough correction for unobserved heterogeneity.[72] I treat each enrollee's future duration enrolled in the market $(nMon_{i,t+k})$ as fixed as observed in the data. Given these assumptions, it is easy to show that:

$$\frac{\partial D_{ij}^{Fut}}{\partial P_j} = \frac{\partial S_{ij}^{Curr}}{\partial P_j} \cdot \left(\sum_{k \geq 1} \rho^k \cdot nMon_{i,t+k} \right) \tag{17}$$

where $\partial S_{ij}^{Curr} / \partial P_j$ is the effect of price on current year's choice share.

Finally, I need to specify insurers' future profit margins. Although imperfect, I simply assume that insurers expect M_{ij}^{Fut} to equal current margins at the enrollee level – which assumes that prices and costs grow in parallel for each enrollee. Note that I still treat M_{ij}^{Fut} as a constant in the pricing FOC (i.e., it does not have a derivative with P_j) but plug in the equilibrium current profit margin ($= \varphi_i P_j^* - c_{ij}\left(N_j\right)$) for it at the end of the calculation.

[72] I base this on the choice persistency of re-enrollees, people who leave the market and make an active plan choice upon their return. Re-enrollees tend to choose the same plan as during their prior spell about 55% of the time. I assume that absent inertia, current enrollees would have stuck with their current plan this frequently. Thus, if the inertia probability is ρ, the total non-switching probability for current enrollees would be $\rho + (1-\rho) \cdot 55\%$. Solving for the value that leads the non-switching probability to equal 95% yields $\rho = 89\%$.

Combining these assumptions and defining the term in parentheses in (17) as $nFutMon_i$, the pricing FOC for insurer j is:

$$
\begin{aligned}
0 &= \frac{\partial \pi_j}{\partial P_j} + \sum_i M_{ij}^{Fut} \cdot \frac{\partial D_{ij}^{Fut}}{\partial P_j} \\
&= \sum_i \varphi_i \cdot nMon_i \cdot S_{ij}^{Curr}(.) + \sum_i \left(\varphi_i P_j - c_{ij}\right)\left(nMon_i + nFutMon_i\right) \cdot \frac{\partial S_{ij}^{Curr}}{\partial P_j}
\end{aligned} \tag{18}
$$

Accounting for future profits adds the $nFutMon_i$ term to the FOC, which increases the incentive to lower prices (just like a steeper demand curve). This effect is likely to have a significant impact. Months in the current year $(nMon_i)$ average 6.2, and future months implied by inertia $(nFutMon_i)$ average 6.8. So the future profit effect works like a more than doubling of the demand slope.

Appendix Figure 1. Hospital Coverage in Massachusetts Exchange Plans

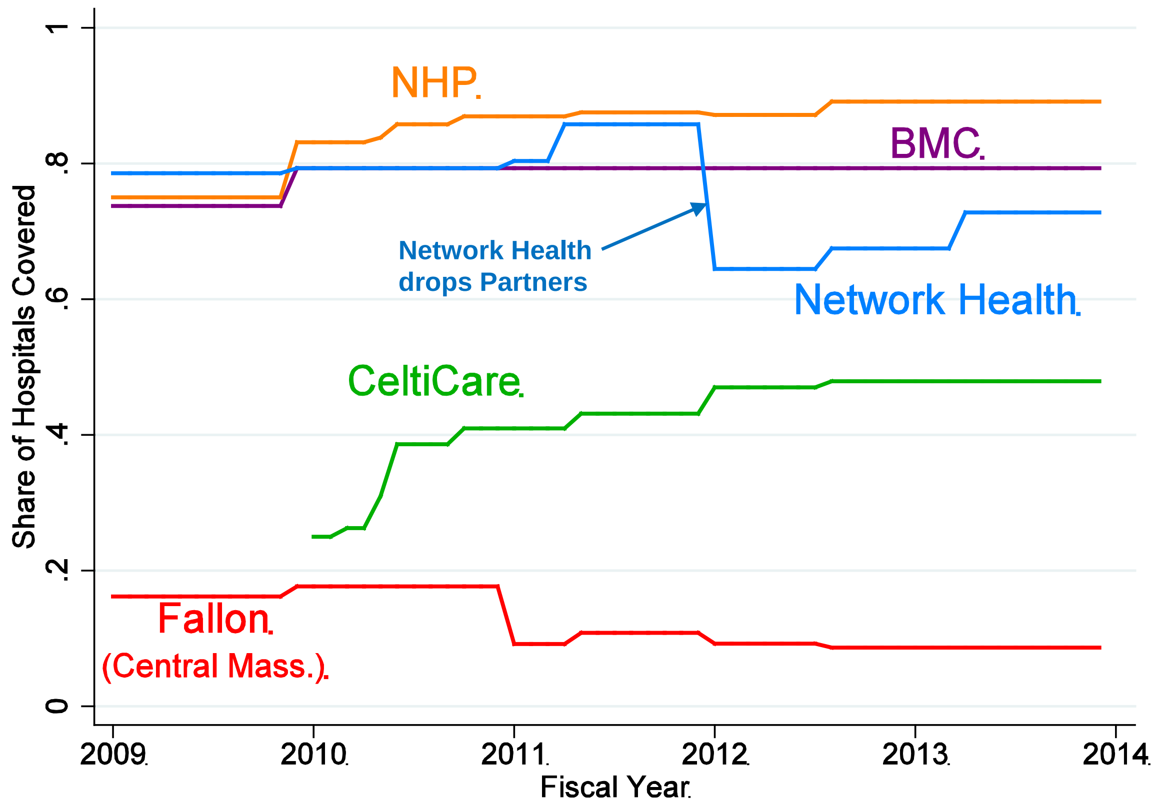

NOTE: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a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

Appendix Table 1. Data Summary Statistics

Hospital Choice Sample

Patient Characteristics

Variable		Mean
No. of Hospitalizations		74,308
Age		44.6
Male		49%
Emergency Department		56%
Diagnoses	Mental Illness	16.7%
	Digestive	13.5%
	Circulatory	12.6%
	Injury / Poisoning	7.1%
	Respiratory	7.0%
	Cancer	6.4%
	Endocrine / Metabolic	6.0%
	Musculoskeletal	5.6%
	Genitourinary	5.1%
	Pregnancy / Childbirth	5.0%
	All Others	14.9%

Chosen Hospital Statistics

Variable	Mean	Std. Dev.
Distance: Chosen Hosp. (miles)	14.1	16.3
All Hospitals (miles)	48.4	25.9
Hospital Category		
Academic Med. Ctr.	29%	---
Teaching Hospital	19%	---
All Others	52%	---
Partners Hospital	14%	---
Out-of-Network	8%	---
Past Use of Chosen Hospital (prior to this year)		
Any Use	42%	---
Inpatient Use	14%	---
Outpatient Use	40%	---
Total Cost to Insurer	$11,369	$15,711
Price (estimated)	$10,981	$4,112
Patient Severity (estimated)	1.000	0.310

Plan Choice Sample

Enrollee Characteristics

Variable	Mean	Std. Dev.
No. of Enrollees	611,455	---
Age	39.6	13.8
Male	46.5%	---
Income: <100% Poverty	47.1%	---
100-200% Poverty	39.6%	---
200-300% Poverty	13.3%	---
Past Hospital User	46.3%	---
Partners Hospitals	7.9%	---
Other Hospitals	42.5%	---
Risk Adjustment Score	0.99	0.90
Choice Type: New Enrollee	29.5%	---
Re-Enrollee	13.5%	---
Current Enrollee	57.1%	---

Plan Choice Statistics

Variable	Mean	Std. Dev.
No. of Choice Instances	1,588,889	---
Insurer Price	$380.7	$69.5
Cons. Premium: Below Poverty	$0.0	$0.0
Above Poverty	$47.3	$45.7
Costs per Month: Total	$371.5	$1,480
Hospital Inpatient	$81.5	$1,048
Non-Inpatient	$290.0	$873
Current Enr: Non-Switching	95.8%	---
Market Shares: BMC	35.5%	---
Network Health	34.7%	---
NHP	19.2%	---
CeltiCare	6.9%	---
Fallon	3.8%	---

Appendix Table 2. Unused Observables Test for Adverse Selection

	Dependent Variable:				
	Hospital Use and Cost (Plans Covering Partners Only)				**Plan Choice**
	Share of Admissions at Star Hospital	Price per Admission ($)	Hospitalization Rate (annual)	Total Cost per Year ($)	Actively Choose Plan Covering Partners
	(1)	(2)	(3)	(4)	(5)
Past Patient at Partners Facility	0.322**	3,143.0**	0.0039	1,137.3**	0.298**
	(0.010)	(126.8)	(0.0034)	(96.3)	(0.004)
Control Variables					
Risk Score	0.007**	99.2**	0.0953**	4,788.7**	-0.004**
	(0.002)	(22.3)	(0.0040)	(159.0)	(0.001)
Plan x Year x Income Grp FE	X	X	X	X	---
Year x Income Grp FE	---	---	---	---	X
Observations	10,505	10,505	270,198	270,198	172,874
R-Squared	0.184	0.127	0.029	0.117	0.181
Dependent Var. Means:					
Past Patient at Partners Facility	0.398	14,125	0.111	7,318	0.661
All Others	0.066	10,770	0.072	4,032	0.376
[Difference]	*[0.332]*	*[3,355]*	*[0.039]*	*[3,286]*	*[0.285]*

** p<0.01, * p<0.05

NOTE: The table shows regression results from the unused observables test for adverse selection, as described in Section 3.1. The bottom section shows raw means of the dependent variables, to show how controlling for risk affects the between-group differences. The data are at the individual x plan choice instance level for the 2011-2013 period during which I have full risk adjustment data. Cost and hospital use outcomes are defined as averages over the subsequent year. “Past Patient at Partners Facility” is a dummy for whether an individual has been observed using a Partners facility for outpatient care prior to the given plan choice instance. The sample excludes new enrollees into the exchange, for whom past utilization data is not observed. Columns (1)-(4) limit the sample to plans covering Partners to examine a sample who all have access to the star hospitals. Column (5) limits the sample to individuals making active plan choices when re-enrolling in the exchange after a gap in coverage. Regressions in columns (4) and (5) are weighted by the number of months each individual was enrolled during the year. All standard errors (in parentheses) are clustered at the individual level.

Appendix Table 3. Robustness Analysis for Unused Observables Test

	Dependent Variable: Total Cost per Year ($)				
	Baseline Regression	Sample: Dx-Based Risk Adj. Only	Sample: Re-Enrollees Only	Past Doctor Visits Only	Other Past Hospital Use Controls
	(1)	(2)	(3)	(4)	(5)
Past Patient at Partners	1,137.3**	908.2**	1,067.0**	1,304.6**	915.4**
	(96.3)	(108.0)	(162.8)	(160.0)	(93.5)
Other Past Use Variables:					
Any Academic Med. Ctr.					237.1**
					(68.9)
Any Hospital					360.6**
					(73.7)
Control Variables					
Risk Score	4,788.7**	4,561.6**	4,585.6**	4,799.4**	4,752.0**
	(159.0)	(171.0)	(302.4)	(159.4)	(162.9)
Plan x Year x Income Grp FE	X	X	X	X	X
Observations	270,198	172,520	71,303	270,198	270,198
R-Squared	0.117	0.132	0.103	0.117	0.117
Dependent Var. Means:					
Past Patient at Partners	7,318	7,216	6,929	8,088	7,318
All Others	4,032	3,966	4,392	4,321	4,032
[Difference]	*[3,286]*	*[3,250]*	*[2,537]*	*[3,767]*	*[3,286]*

** p<0.01, * p<0.05

NOTE: The table shows robustness checks on the cost regression results in . See the note to that table for additional descriptions of the data and regression setup. Column (1) replicates the cost regression in column (4) of Appendix Table 2. The remaining columns adjust the sample or regression variables relative to this baseline. Column (2) restricts the sample to individuals for whom the best quality, diagnosis-based risk adjustment is available. Column (3) restricts the sample to re-enrollees only, who are not subject to inertia because (unlike current enrollees) they must make an active plan choice. Column (4) uses the full sample but defines past Partners patients only based on past doctor visits, not other forms of outpatient care. Column (5) controls for additional past use covariates, including past use of any hospital and past use of any academic medical center.

Appendix Table 4. Analysis of Costs for Network Health Enrollees, 2011-12

Enrollee Group	Avg. Costs			Risk-Adjusted Avg. Costs			Group
	2011	2012	%Δ	2011	2012	%Δ	Size*
All Enrollees	**$4,631**	**$3,676**	**-21%**	**$4,439**	**$3,761**	**-15%**	**---**
Stayers	$3,877	$3,641	-6%	$3,807	$3,596	-6%	36,768
Left Plan in 2012							
Switched Plans	$8,045	*[$7,391]*	---	$6,109	*[$5,106]*	---	4,640
Exited Market	$5,634	---	---	$5,511	---	---	22,617
Joined Plan in 2012							
Switched Plans	*[$3,391]*	$3,461	---	*[$3,641]*	$3,706	---	15,062
Entered Market	---	$3,781	---	---	$4,007	---	51,109

Breakdown by Partners Patient Status

Enrollee Group	Avg. Costs			Risk-Adjusted Avg. Costs			Group
	2011	2012	%Δ	2011	2012	%Δ	Size*
Stayers							
Partners Patients	$6,228	$4,633	-26%	$5,533	$3,933	-29%	5,217
Other Dropped Hosp.	$4,952	$4,354	-12%	$4,089	$3,378	-17%	3,643
All Others	$3,227	$3,382	5%	$3,334	$3,562	7%	27,908
Switched from Network Health in 2012							
Partners Patients	$9,249	*[$8,056]*	---	$6,802	*[$5,264]*	---	3,160
Other Dropped Hosp.	$5,723	*[$6,539]*	---	$4,331	*[$4,609]*	---	825
All Others	$4,043	*[$5,197]*	---	$3,910	*[$4,839]*	---	655
Exited Market in 2012							
Partners Patients	$9,871	---	---	$7,498	---	---	3,104
Other Dropped Hosp.	$7,431	---	---	$5,490	---	---	1,659
All Others	$4,716	---	---	$5,019	---	---	17,854

* Number of enrollees during the relevant year they were enrolled in Network Health.

NOTE: The top panel of this table shows the changes in medical costs per member-month for Network Health from 2011 (when it covered the star Partners hospitals) to 2012 (when it dropped them). The first set of columns show raw, unadjusted costs. The next columns show risk-adjusted costs, defined as a group's average cost divided by its average risk score. Group size is the number of enrollees in the relevant group during the year(s) they were enrolled in Network Health. Overall, Network Health's costs fell by 21%, or 15% after risk adjustment. The next rows break costs into enrollee subgroups: a fixed group of "stayers" (people in the plan in both years) and enrollees who left or newly joined the plan in 2012. The bottom panel breaks down the results for stayers, switchers, and exiters between past Partners patients (prior to the start of 2012), past patients of other hospitals dropped by Network Health in 2012, and all others.

Appendix Figure 2. Plan Switching and Selection for Network Health: Below-Poverty Enrollees Only

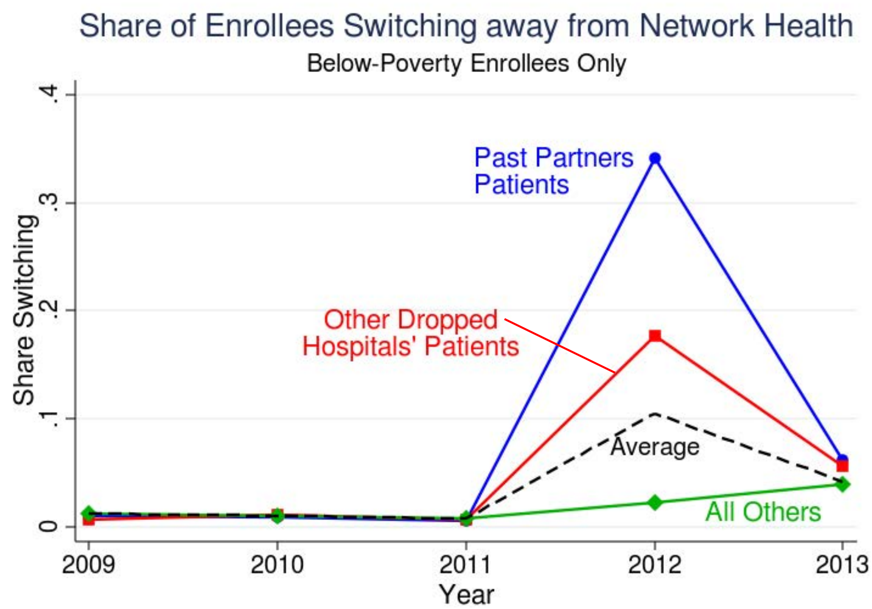

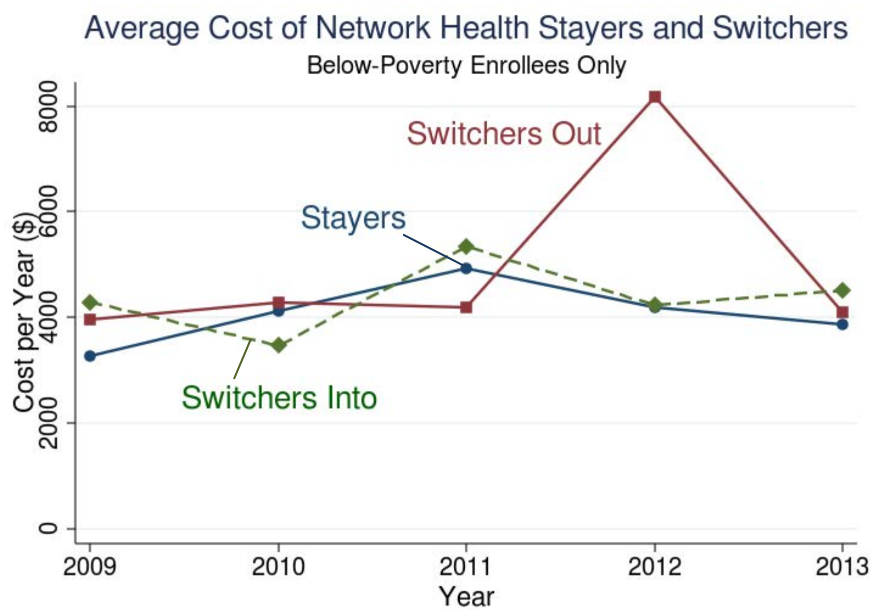

NOTE: These figures are exactly analogous to Figure 3 but with the sample restricted to below-poverty enrollees only, for whom there are no premium changes (all plans are free in all years). See the note to Figure 3 for additional details. Both the switching patterns away from Network Health and the average cost of switchers and stayers are qualitatively similar for this below-poverty group as for the full sample, suggesting that the results are driven by the network change not by concurrent premium changes.

Appendix Figure 3. Model Fit for Plan Average Medical Costs

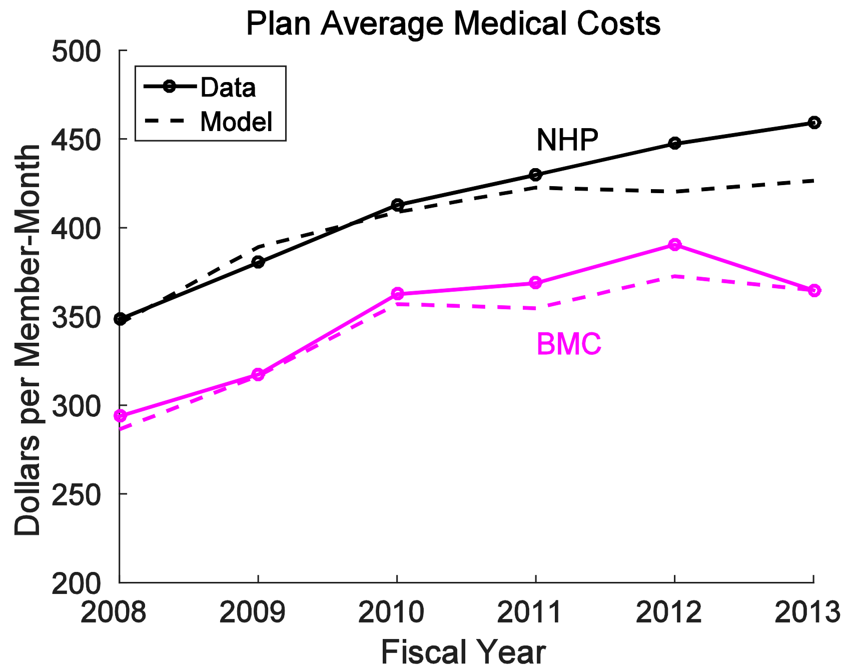

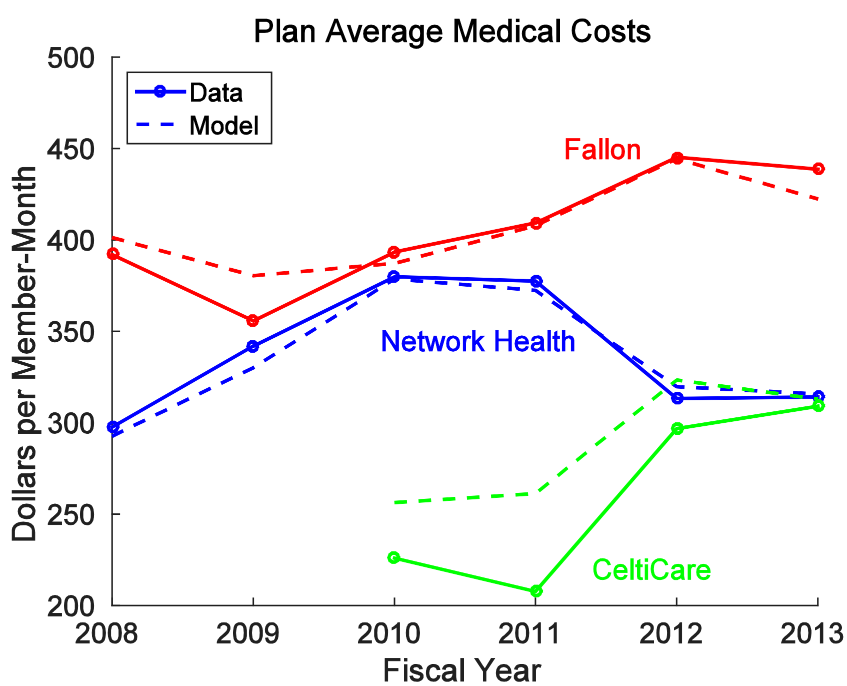

Appendix Figure 4. Plan Switching Patterns

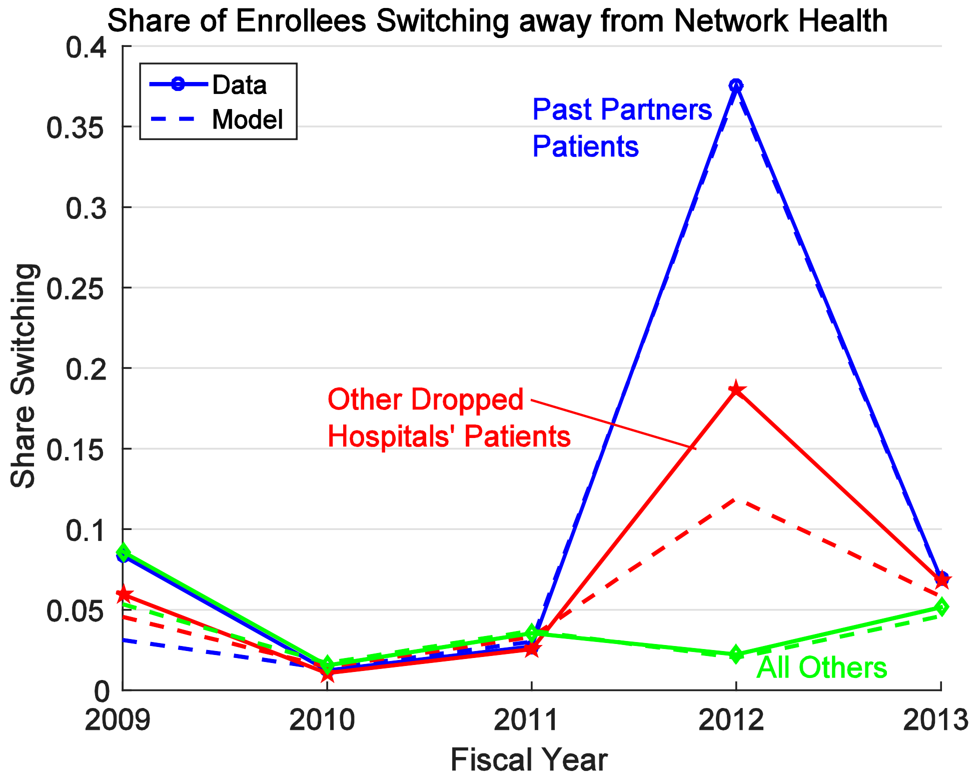

Appendix Table 5. Cost Changes for Network Health

Network Health: Average Costs 2011-12

Enrollee Group	**Data**				**Model**			
	2011	**2012**	**%Δ**	**Risk Adj. %Δ**	**2011**	**2012**	**%Δ**	**Risk Adj. %Δ**
All Enrollees	**$378**	**$313**	**-17%**	**-15%**	**$372**	**$320**	**-14%**	**-14%**
Stayers (in plan both years)	$317	$305	-4%	-5%	$331	$311	-6%	-8%
2011 Only Enrollees	$476	---	---		$439	---	---	
2012 Only Enrollees	---	$310	---		---	$320	---	

Appendix Figure 5. Admission Shares at Hospitals Dropped by Network Health in 2012

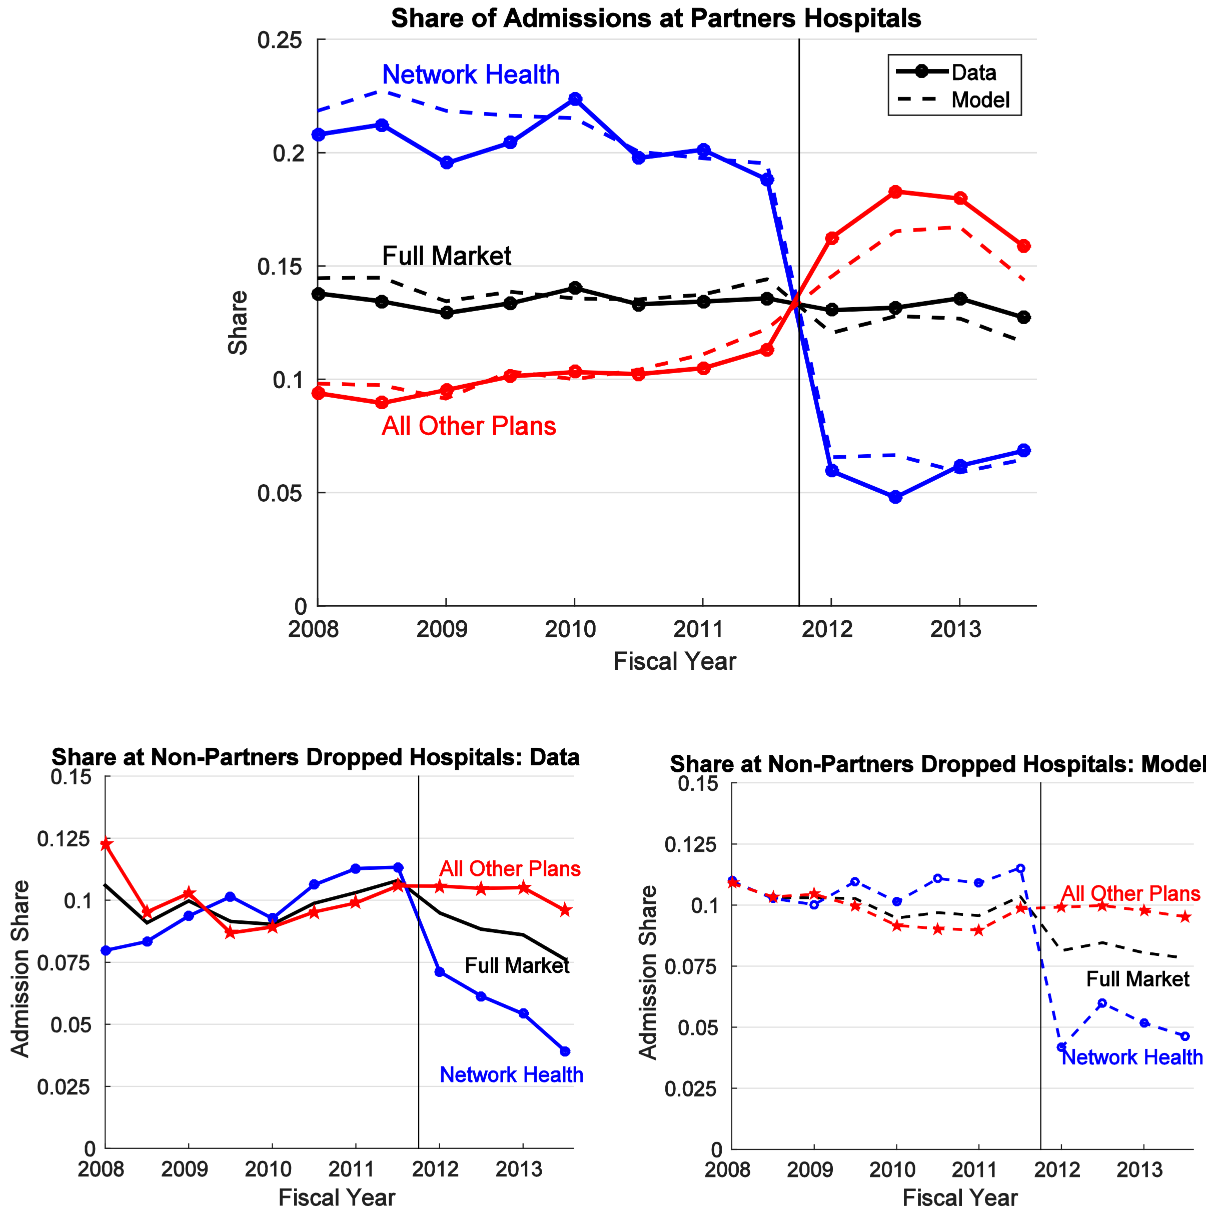

NOTE: These figures show the share of hospital admissions at hospitals that Network Health plan dropped from its network in 2012. The dashed lines show the model's prediction for the same statistics. These are calculated holding fixed each individual's observed plan, not reassigning plan choices using the plan demand model.

Appendix Figure 6. Changes in Cost per Hospital Admission around 2012 Network Changes

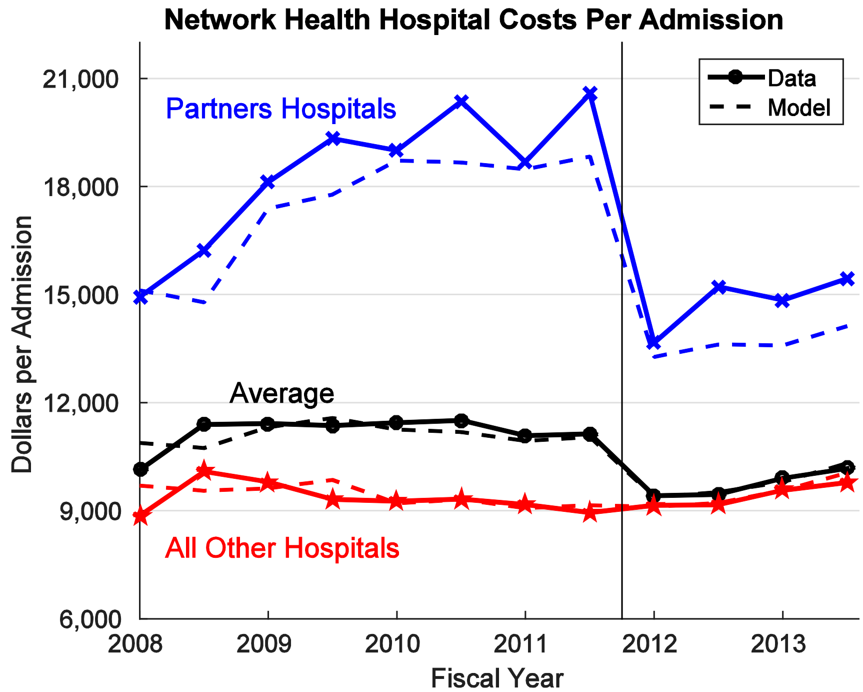

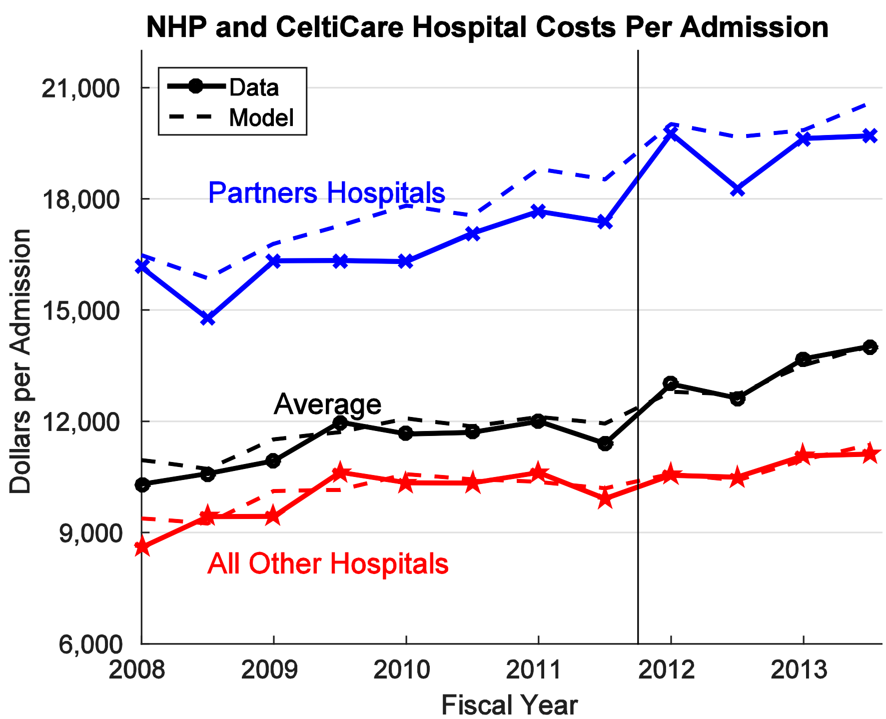

NOTE: These figures show average costs per hospital admission for two sets of plans: Network Health (top figure), which dropped the star Partners hospitals in 2012, and NHP and CeltiCare (bottom figure), which continued to cover them. The dashed lines show the model's prediction for the same statistics. These are calculated holding fixed each individual's observed plan, not reassigning plan choices using the plan demand model.