

Health care demand elasticities by type of service

Randall P. Ellis, Bruno Martins, and Wenjia Zhu

Boston University
Department of Economics
270 Bay State Road
Boston, Massachusetts 02215
USA

ellisrp@bu.edu (corresponding author), bmartins@bu.edu, wz97@bu.edu

October 26, 2016

Abstract

We estimate within-year price elasticities of demand for detailed health care services using an instrumental variable strategy, in which individual monthly cost shares are instrumented by employer-year-plan-month average cost shares. Using 172 million person-months spanning 73 employers from 2008-2014, we estimate that the overall elasticity by myopic consumers is -0.33, with higher elasticities for pharmaceuticals (-0.33), specialty care (-0.24), and mental health/substance abuse (-0.18), and lower elasticities for prevention visits (-0.01), and emergency rooms (-0.03). Overall demand is more elastic at higher cost shares, in single than in family coverage plans, among older, healthier people than for younger, sicker individuals.

Keywords: Elasticity; cost sharing; high deductible health plans; dynamic demand for health care; myopic expectations.

(JEL: I11, C21, D12)

Acknowledgements: Research for this paper was supported by the National Institute of Mental Health (R01 MH094290). The views expressed here are the authors' own and not necessarily those of the NIMH. We are grateful to Zarek Brot-Goldberg, Tom McGuire and attendees at the University of Chicago (Will Manning Memorial Conference), Boston University, the University of Oxford, and the American Society of Health Economists for comments on previous versions.

1. Introduction

This paper revisits the classic issue of price elasticities of demand for health care services, a topic of central importance for understanding moral hazard and optimal insurance plan design, as well as having important implications for health plan choice, access to care, health care cost containment, risk adjustment, and financial risk. Moreover, since modern insurance plans incorporate highly specific benefit plan features in their designs, and these features can be customized for demographic, employment, or geographically-defined groups, understanding demand responses for disaggregated service types is of significant policy interest. In this paper we use a novel instrumental variable technique on a very large sample (172 million person-months, 73 employers, 1468 health plans) to generate estimates of short-run own- and cross-price elasticities of demand among the privately insured for both aggregated and highly disaggregated types of services and population subgroups. Many of our estimates, particularly those of services that are not widely used (e.g., MRIs, mental health/substance abuse treatment, lab tests, and emergency department use), are infeasible to make on small to modest size samples.

In estimating price elasticities of demand for health care, a notable challenge is to control for endogeneity of prices stemming from endogenous health plan choice. Several estimation strategies have been used in previous studies to overcome this difficulty, including randomized control trials (RCTs) (Manning et al. 1987), natural experiments (Duarte 2012; Brot-Goldberg et al. 2015), within-year price variation arising from nonlinear health insurance contracts (Ellis 1986; Ellis and McGuire 1986), and instrumental variable strategies (Eichner 1998; Einav et al. 2015; Kowalski 2016; Scoggins and Weinber 2016).

The classic study by Manning et al. (1987) estimated health care demand responses using data from the RAND Health Insurance Experiment (HIE) RCT conducted in the mid-1970s. The HIE randomly assigned around 6000 individuals to fourteen health plans with cost shares that ranged zero to 95 percent and stoplosses (limits on out of pocket payments) that ranged from zero to \$1000, enabling relatively precise estimates of demand response, in aggregate and for certain broad categories of services and subpopulations. Aron-Dine et al. (2013) reanalyzed the HIE data while highlighting the importance of incorporating the nonlinearity features of insurance contracts that is common today in the estimation of price elasticity of demand for health care. They show that the Manning et al. (1987) demand elasticity estimate of -0.2 is lower

than other calculations would suggest, estimating an overall demand elasticity of -0.5 once free cost share plans are excluded.

Also using the HIE data was Keeler and colleagues (1988a, 1988b) who analyzed episodes of health care treatment. This refinement allows them to evaluate how insurance affects the probability of seeking treatment and conditional spending if treatment is sought, as well as enables them to model expectations about cost shares more carefully than previous studies. Zweifel and Manning (2000) demonstrate that relying on simple cross-sectional variation across health plans would contaminate demand response estimates due to endogenous selection into health plans: less generous (higher cost sharing plans) will on average attract a different and usually healthier group of enrollees than more generous (lower cost sharing) health plans.¹

Since RCTs such as the HIE are expensive and rare, many studies have exploited natural experiments that induce changes in prices to estimate demand responsiveness. Duarte (2012) is a good recent example: he uses Chilean data with a large number of diverse health plans to study the demand response of both elective (home visits, psychologist visits, physical therapy evaluations) and acute (appendectomy, cholecystectomy, arm cast) health care services when employers change their coverage generosity. In the similar spirit, Brot-Goldberg et al. (2015) exploit a natural experiment in health insurance coverage due to a required switch to a high-deductible plan by one large company. They estimate that switching from free care to a high-deductible plan covering about 80 percent of total health expenditures reduced overall spending by 11.8% to 13.8%, driven mostly by reductions in quantity of services. They find a considerable demand response even among sick people who expect to exceed deductibles and face low end-of-year prices, suggesting that people are myopic, responding strongly to the spot prices at the time of care is received as much as to the more theoretically correct expected end-of-year price.

Another estimation approach is to take advantage of the within-year price variability that occurs because of deductibles, coverage ceilings, and out-of-pocket stoplosses. Building on the HIE work of Keeler, Manning and colleagues, Ellis (1986) and Ellis and McGuire (1986) were

¹ Cutler et al. (2008) show that empirically the bias can either be favorable or adverse, so this statement is not always correct.

among the first to estimate demand response while relying solely on this variation, studying both within-year demand responsiveness of mental health services to a coverage ceiling, and demand response to deductibles at a single large firm that increased its deductibles and stoplosses. Other studies using the nonlinear price schedule include Meyerhoefer and Zuvekas (2010) and Meyerhoefer et al. (2014). Each of these studies uses individual- or household-level variables to control for biased plan enrollment, but leaves selection due to within-year price variation uncontrolled. This can introduce bias, since people who exceed a deductible will on average be sicker than those who do not, and hence one would expect them to consume more services for this reason alone, not only because of the change in coverage.

A final approach is to use instruments to control for the endogeneity of prices (cost shares) within a given year. Eichner (1998) proposes a novel instrument for within-year price variation: the occurrence of accidental injuries (e.g., a broken leg) by dependents as a convincingly exogenous event that changes subsequent cost sharing and hence affects the employee's own demand for medical care. Building on the Eichner framework, Kowalski (2016) uses emergency department spending for injuries by other family members in households with at least four members for her estimation of demand responsiveness, where the instrument (whether a family member incurs spending for an injury) derives its power from the fact that the event increases the likelihood that the family deductible or stop loss can be exceeded, thus reducing the marginal price of care faced by the remaining household members. Einav et al. (2015) use changes in the Medicare Part D coverage for prescription drugs as an instrument to study consumer responses to cost sharing, while Scoggins and Weinber (2016) use the plan characteristics as an instrument for the actual observed cost share in claims data. A common feature of all of these studies is that only one (or a small number of) distinct health plans are available for the study, so they rely on individual (or household) level instruments for identification. The challenge is that individual level attributes are often correlated with demand for care and one must be very careful about selecting variables used as instruments in order not to violate the exclusion restriction. Moreover, elasticities based on a limited number of plans might not adequately meet external validity, as those plans studied may not be representative of the entire population.

The contribution of this paper is twofold. First, we model the effects of within-year cost sharing variation using aggregate rather than individual level instruments for identification. In our sample, each person contributes twelve monthly observations on endogenous cost sharing

and spending decisions. To control for this cost sharing endogeneity, we use the employer-, year-, and plan-level average monthly cost share as an instrument for individual cost shares, exploiting the fact that different firms, and different plans have widely different within-year variation in their cost sharing. We provide evidence that these employer average cost shares at the plan level are very weakly correlated with the average health status of plan enrollees. We are not aware of any other research studies that have had a sufficiently large sample of employers and health plans as to enable employer-, year-, and plan-level average monthly cost shares to be used as instruments. Moreover, our instrument does not rely on any detailed characteristics of health plans, such as the actual value of deductibles and coinsurance rates, which makes it a useful instrument under limited plan data. Our econometric models further control for individual*year fixed effects to absorb not only individual, but plan, employer, county, and year fixed effects, so the only within-year variation of interest is cost sharing variation.

Second, we use highly detailed micro data on 172 million person-months, which allows us to estimate elasticities for very detailed types of health care services, which is usually not possible due to low utilization rates in such services. Understanding how consumers react to change in cost shares for different services is important to understand how insurers design their health plans to control for moral hazard and to select services to offer. In that sense, our paper is similar to Einav et al. (2016), which use detailed micro data to estimate elasticities for more than 150 pharmaceutical drugs in Medicare Part D and analyze to the extent cost share and elasticities are correlated. In Ellis et al. (2016), our estimated elasticities are used to evaluate health plan incentives for service-level selection.

We use commercial claims and encounter data from 2007 to 2014 to estimate demand responsiveness and elasticities. Our estimation method uses regression techniques to estimate unconditional spending responsiveness, using employer level averages as instruments for the individual own cost share. We control for individual*year and monthly time fixed effects. Patient health, geography, employer-year and plan characteristics are controlled for by individual*year fixed effects, while endogeneity of cost share (the consumer's price) is controlled for by IV techniques.

As a precursor of our findings, we find clear evidence of short run elasticities of demand in response to within-year variations in cost sharing: people do respond to within-year variation. For the overall demand elasticity, assuming myopic expectations yields an estimate of -0.33

while assuming forward expectations yields an estimate of -0.31. These estimates are slightly larger than those in the Rand HIE which were approximately -0.2, but lower than the elasticities estimated by Aron-Dine et al. (2013) and Kowalski (2016), which we suggest below is because in those studies the average consumer cost shares were higher than typical among private firms. We also find heterogeneity in the price elasticity of demand across different types of services, and across health plans, with higher elasticities for pharmaceuticals (-0.33), specialty care (-0.24), and spending on mental health/substance abuse (-0.18), and lower elasticities for prevention visits (-0.01), emergency rooms (-0.03), and mammogram (-0.07). Supporting Einav et al. (2013), we find evidence of not only biased selection but also heterogeneity in price responsiveness, with consumers in health plans that have large changes in cost sharing (e.g., HDHP) yielding higher estimated demand elasticities than HMOs. Demand elasticities also vary by enrollee age, risk scores, time, salary versus wage earners, and industry.

2. Data

We use eight years of Truven Analytic's MarketScan commercial claims and encounter database between 2007 and 2014. This claims data source was also used to develop and calibrate the Accountable Care Act Health Insurance Marketplace risk adjustment formulas, and has broadly distributed samples of enrollees primarily at large employers. Further steps used in cleaning up the data follow the steps described in Ellis and Zhu (2016). Our estimation uses a sample of 73 large employers with enrollees present in the sample frame for all twelve months for at least three years. We sum up total covered payments and out-of-pocket (OOP) payments in each calendar month for various types of services, and used this information to calculate the actual cost share paid (the ratio of OOP to covered amounts) in each month for each service.² All dollar amounts were converted into December, 2014 dollars by dividing by the appropriate monthly Personal Consumption Expenditure (PCE) deflator for health care costs from the US Bureau of Economic Analysis. Spending was also adjusted for the number of days in the month

² We assigned zero monthly spending if the spending for that month, in each type of service, was lower than 1USD, as these values are likely to be adjustments and claim-processing errors. If the total out-of-pocket spending for a month was negative, we changed it to zero for the same reason.

to remove this minor monthly variation in spending. Hence remaining changes over monthly spending reflect primarily the gradual effects of the aging, technological change, and health plan features rather than inflation or uncontrolled seasonality. The effects of this aging and gradual technological change are eliminated by using monthly time fixed effects and including plans with no cost sharing change within a year. Hence in our sample, changes in spending across months only reflect the effects of changes in cost sharing within the year.

For observations in which monthly spending by an individual is positive, we define cost share as the ratio between out-of-pocket and spending. If monthly spending is zero we cannot compute cost-share, nor do we know the out-of-pocket payment that the individual was expecting to pay if she actually had positive utilization. Knowing the value of the counterfactual cost share if the individual had positive spending is of considerable interest for studying the decision to seek care, yet there is little information about how such expectations are formed. We use two types of individual expectations that we hope will bracket likely price expectations: a myopic expected price and a forward-looking expected price.

We model cost sharing impacts within a year at the health plan level, where for this paper we define a “health plan” as defined by a unique plan identifier and whether the individual chooses family or single coverage. Our estimation sample has available data from 73 employers with 1468 health plan years and information on 14.3 million individual years. This large panel enables us to estimate own-price demand elasticities not only for common, but also for relatively rare disaggregated health care services. We parse total spending in three ways, first into three broad categories of service (outpatient, inpatient, and prescription drugs), second into 11 finer types of service, and finally into 33 disaggregated types of service. This finest classification is closer to the level at which US health plans and employers are able to adjust coverage generosity in order to influence plan selection and costs.

We processed the claims of all individuals in our sample using the Verisk Health’s DxCG Medical Intelligence software, which generated DxCG risk scores that predict the next year’s health care total spending. These risk scores, which average one in their original estimation sample, express how expensive a given individual is expected to cost relative to the population average. Hence a person expected to cost twice the population average would have a relative risk score (RRS) of 2.00. Using a prospective risk score requires that we drop the first year of spending and cost share information for each unique individual in our sample.

Table 1 shows key sample characteristics. Across all years combined, we have a sample of around 14.3 million individual years with an average age of 34.1 years.

Table 1
Characteristics of individuals in the sample

Year	Obs.	Individuals	Age	Female	Concurrent risk score	Prospective risk score
2008	27,227,964	2,268,997	34.16	0.52	1.00	0.99
2009	26,633,708	2,219,476	34.03	0.52	1.11	1.02
2010	30,290,976	2,524,248	34.43	0.52	1.16	1.10
2011	29,591,892	2,465,991	34.02	0.51	1.15	1.09
2012	25,271,124	2,105,927	33.76	0.50	1.14	1.07
2013	15,744,564	1,312,047	34.56	0.51	1.28	1.10
2014	17,361,960	1,446,830	34.07	0.50	1.33	1.07
ALL	172,122,188	14,343,516	34.13	0.51	1.15	1.06

2.1. Types of services modeled

As is well known, aggregation across health care services may introduce spurious estimates of the relationship between prices and spending if services with higher spending are also those with better (or worse) coverage (Newhouse et al. 1979). Therefore in addition to modeling total spending, we look at spending by types of service. We aggregate services into meaningful categories in three different ways that differ in their degree of fineness. Our first and coarsest measure is a three-way division between inpatient care, outpatient care and pharmaceutical claims, with services in each category as defined by MarketScan, which means that many inpatient physician services are in fact categorized into outpatient care rather than inpatient care. Our second definition of types of service is based on MarketScan's definitions of service location (surgery ward, mental health facility, etc.) and procedures (MRI, specialty visit, etc.). Using a combination of the two dimensions, we map 576 different combinations into 11 broad categories which we call Aggregate Types of Service (ATOS). Finally, we further disaggregate these 11 types of services into 33 Types of Services (TOS). Table 2 and 3 display summary statistics for the three groupings of categories of service.

Table 2

Monthly summary statistics on the study variables, for two levels of aggregation

Types of service	Mean monthly spending (\$)	% Total spending	Std. dev	Mean cost share	% obs. positive spending	Mean spending if positive
All spending	387	100.00%	2967	31.01%	46.35%	836
Inpatient	222	57.25%	1396	27.28%	32.15%	690
Outpatient	80	20.69%	2408	8.58%	0.42%	19,089
Rx	85	22.06%	519	40.73%	34.23%	250
Broad categories of service:						
1 - Primary care	32	8.15%	249	25.40%	18.85%	168
2 - Maternity	12	3.21%	433	18.60%	0.39%	3,232
3 - MH/SA	11	2.83%	302	29.74%	3.00%	366
4 - Surgery	13	3.29%	786	4.28%	0.07%	17,512
5 - Specialty care	85	21.96%	918	30.99%	12.39%	687
6 - ER	19	4.96%	289	25.03%	1.42%	1,352
7 - Room & board	32	8.14%	1496	9.46%	0.26%	11,906
8 - Imaging	28	7.24%	392	19.60%	5.78%	486
9 - Lab tests	27	6.95%	341	23.39%	12.75%	211
10 - Pharmacy	112	28.79%	941	39.43%	35.04%	318
11 - Other	18	4.60%	392	20.50%	4.21%	424

Table 3

Monthly summary statistics on the study variables, by detailed Types of Service (TOS)

Type of service	Mean monthly spending(\$)	% Total spending	Std. dev.	Mean cost share	% obs. positive spending	Mean spending if positive
1 - Non-specialty visits	26	6.66%	220	28.14%	17.15%	150
2 - Home visits	1	0.26%	108	7.55%	0.08%	1,219
3 – Prevention	5	1.23%	27	4.35%	3.48%	137
4 – Maternity	12	3.21%	433	18.60%	0.39%	3,232
5 - Mental health	9	2.37%	247	29.87%	2.92%	315
6 - Substance abuse	2	0.46%	170	24.77%	0.10%	1,727
7 - Surgical procedures	8	1.98%	539	4.03%	0.07%	11,549
8 - Surgical supplies/devices	5	1.31%	406	4.83%	0.06%	7,978
9 - Nonsurgery IP procs	8	2.00%	343	10.16%	0.24%	3,181
10 - Specialty visits	63	16.36%	774	30.81%	9.59%	661
11 – Dialysis	3	0.70%	253	4.12%	0.03%	8,322
12 - PT, OT, speech therapy	9	2.29%	114	24.33%	2.95%	301
13 - Chiropractic	2	0.61%	21	37.79%	2.48%	95
14 – Hospice	0	0.01%	17	2.80%	0.00%	3,316
15 - Emergency room	19	4.96%	289	25.03%	1.42%	1,352
16 - Room board - Surgical	20	5.05%	1290	8.26%	0.13%	15,470
17 - Room board - Medical, other	12	3.08%	713	10.22%	0.15%	8,151
18 - CAT scan	6	1.45%	130	18.35%	0.60%	934
19 - Mammograms	3	0.71%	33	6.12%	1.13%	241
20 – MRIS	7	1.70%	124	18.35%	0.56%	1,177
21 - PET scans	1	0.15%	46	7.26%	0.02%	2,471
22 - Radiology - diagnostic	6	1.65%	108	22.55%	3.41%	188
23 - Radiology - therapeutic	3	0.83%	290	4.09%	0.03%	10,722
24 - Ultrasounds	3	0.75%	42	22.71%	1.12%	260
25 - Diagnostic services	11	2.72%	235	17.11%	2.85%	370
26 - Laboratory other	16	4.23%	202	24.04%	11.30%	145
27 – Pharmacy	85	22.06%	519	40.73%	34.23%	250
28 - Facility-based pharmacy	14	3.58%	615	17.07%	0.69%	2,011
29 - Specialty drugs, injections	12	3.15%	429	19.88%	2.60%	470
30 - Non-surg supplies/devices	6	1.44%	210	22.69%	1.54%	363
31 – DME	2	0.64%	89	22.01%	0.69%	361
32 - Transportation	2	0.53%	165	15.23%	0.14%	1,463
33 – Other	8	2.00%	259	17.77%	2.33%	332

2.2. Health plan types

Grouping health plans into plan types is based on the classification used by Truven Health Analytics MarketScan. For this study we classify plans into one of the following plan types: exclusive provider organization (EPO), health maintenance organization (HMO), point of service (POS), preferred provider organization (PPO), comprehensive (COMP), consumer-directed health plan (CDHP) and high deductible health plan (HDHP). MarketScan differentiates the CDHP from HDHP not by the size of the deductibles but instead by whether there is a flexible spending account (CDHP) or a health savings account (HDHP). Some PPOs or POS can also have high deductibles, while not all CDHP have high cost sharing. As Ellis and Zhu (2016) note, plans differ not only in their cost sharing but also in their degree of provider choice, with the HDHP and CDHP having the most choice of providers, and EPOs and HMOs typically having the least. As long as provider choice and other plan variation does not vary systematically within a year, our approach relying on within-year cost sharing variation should be valid.

3. Conceptual framework

We conceptualize consumers as making choices about how much to spend on each health care service in each month primarily in response to shocks, θ_t , that occur at the beginning of the month. Recent studies suggest that consumers are more nearly myopic than having perfect foresight (See especially Einav et al. 2015). Since much of their demand is driven by uncertain developments captured by θ_t , consumers tend to react to new information by making choices about treatment using an expected demand price p_{st}^d for that month for that service, which is the product of the supply price p_s^s and the consumer's actual cost share for service s in month t , CS_{st} which reflects the benefit coverage for that service. Hence we have $p_{st}^d = CS_{st}p_s^s$, and can write the demand for the vector of services $Q_t = \{q_{st}\}_s$ as

$$\begin{aligned} Q_t &= Q[\{p_{st}^d\}|\theta_t] \\ &= Q[\{CS_{st} p_s^s\}|\theta_t] \\ &= Q[\{CS_{st}\}|\{p_s^s\}, \theta_t] \end{aligned}$$

The work of Manning et al. (1987), as well as the more recent work by Aron-Dine et al. (2013) suggests that demand for health care spending is nonlinear, with low levels of cost

sharing having a larger impact on spending per unit than higher levels of cost sharing. As used in Ellis (1986), a demand curve specification that captures this is:

$$q_{st} = \alpha(\theta_{st}) \prod_s \beta_s^{CS_{st}}$$

Or equivalently:

$$\log(q_{st}) = \log(\alpha_{st}) + \sum_s \log(\beta_s) CS_{st} + \varepsilon_{st}$$

This log linear specification is convenient in that the price responsiveness parameters β_s are the exponents of the coefficients on the cost share terms CS_{st} . With sufficient variation in cost shares across services, both the own- and cross-price elasticities can in principle be estimated. In the current paper we focus on own-price demand functions although we also examine specifications that estimate cross-price effects as well.

3.1. Estimation equations

We assume health care spending reflects attributes of the individual (i) (which also uniquely defines a household contract), employer (e), whether the plan is for a family or single coverage (f), health plan including its own unique provider network (p), geographic location such as the county or MSA (c), year (y) and a monthly time trend index (t). Using Y for the dependent variable, and CS^* to denote the consumer's expected cost share (discussed at length below) relevant for the choice of Y , one strategy would be to directly estimate the following.³

$$Y_{it} = \beta CS_{it} + \alpha_i + \gamma_{efpy} + \delta_c + \lambda_y + \gamma_t + \varepsilon_{it} \quad (1)$$

For this paper we estimate instead the following much simpler specification.

$$Y_{it} = \alpha'_{iy} + \beta CS_{it} + \gamma_t + \varepsilon_{it} \quad (2)$$

where $\alpha'_{iy} = \alpha_i + \gamma_{efpy} + \delta_c + \lambda_y$

³ In a working paper, Luo et al. (2016) develop an estimation method capable of estimating equation (1) directly in very large samples.

This equation estimates the relation between spending and cost share controlling only for individual-year fixed effects and time trend (year-month) fixed-effects, which has the advantage of speeding up computation. Since including individual-year fixed effects implicitly controls for individual characteristics that do not vary within the year, we are, in fact, controlling for factors such as age, gender, risk scores, MSA, plans and any other variables that are fixed within a year. The monthly dummies pick up seasonality, changes in national trends, and the effects of any shocks to the economy (such as recessions). The only covariate over time (in months) is the individual's monthly expected cost share.

As discussed below, we estimate an unconditional spending model using the natural log of spending, where we add one to accommodate months with zero spending, the preferred specification of Aron-Dine et al. (2013) using annual data. A log linear specification has the added advantage of not only compressing the effect of the high cost upper tail, but also of estimating a nonlinear demand curve, such that the proportional reduction in demand as the cost share is increased by a fixed amount, such as 0.05, becomes progressively smaller.^{4,5}

Since our model is a log-linear specification, we obtain elasticities by multiplying the estimated price coefficient by the average cost share for that particular service.

3.2. A graphical approach

Our identification relies on variation over time in average cost shares as well as variation across employer*year*family/single*health plans in the benefit coverage of the plans offered. The former variation is required to identify the effect of consumer price expectations on demand, while the second variation is required for our instrumental variable strategy to work (our approach would not work if there is only one employer*plan*year). Figure 1 shows patterns of

⁴ An alternative to unconditional spending models is the two-part model, in which the first stage models the binary choice of seeking care and the second stage models the continuous choice of how much treatment to obtain. While useful to disentangle the demand responsiveness into an extensive margin and an intensive margin, we do not examine the two-part model here as our primary goal is to estimate price elasticities of spending for all services and by detailed types of services.

⁵ Our estimation program, written in SAS, uses 172 million observations, absorbs 14.3 million person*year fixed effects, does two stage least squares, and performs robust clustered error corrections with 16,886 clusters in 68 minutes.

average monthly costs by plan type. It shows the considerable variation across months, across plan types, and across broad types of service. There is even more heterogeneity in our data in that cost shares vary across employers and over time, which is averaged out in these diagrams. High deductible plans have the steepest decline in their average cost shares overall, while HMOs have the least, and have costs shares that are essentially flat. CDHPs, comprehensive plans, PPOs, POS, and EPOs lie intermediate between the two. Comparing the four panels in the figure,

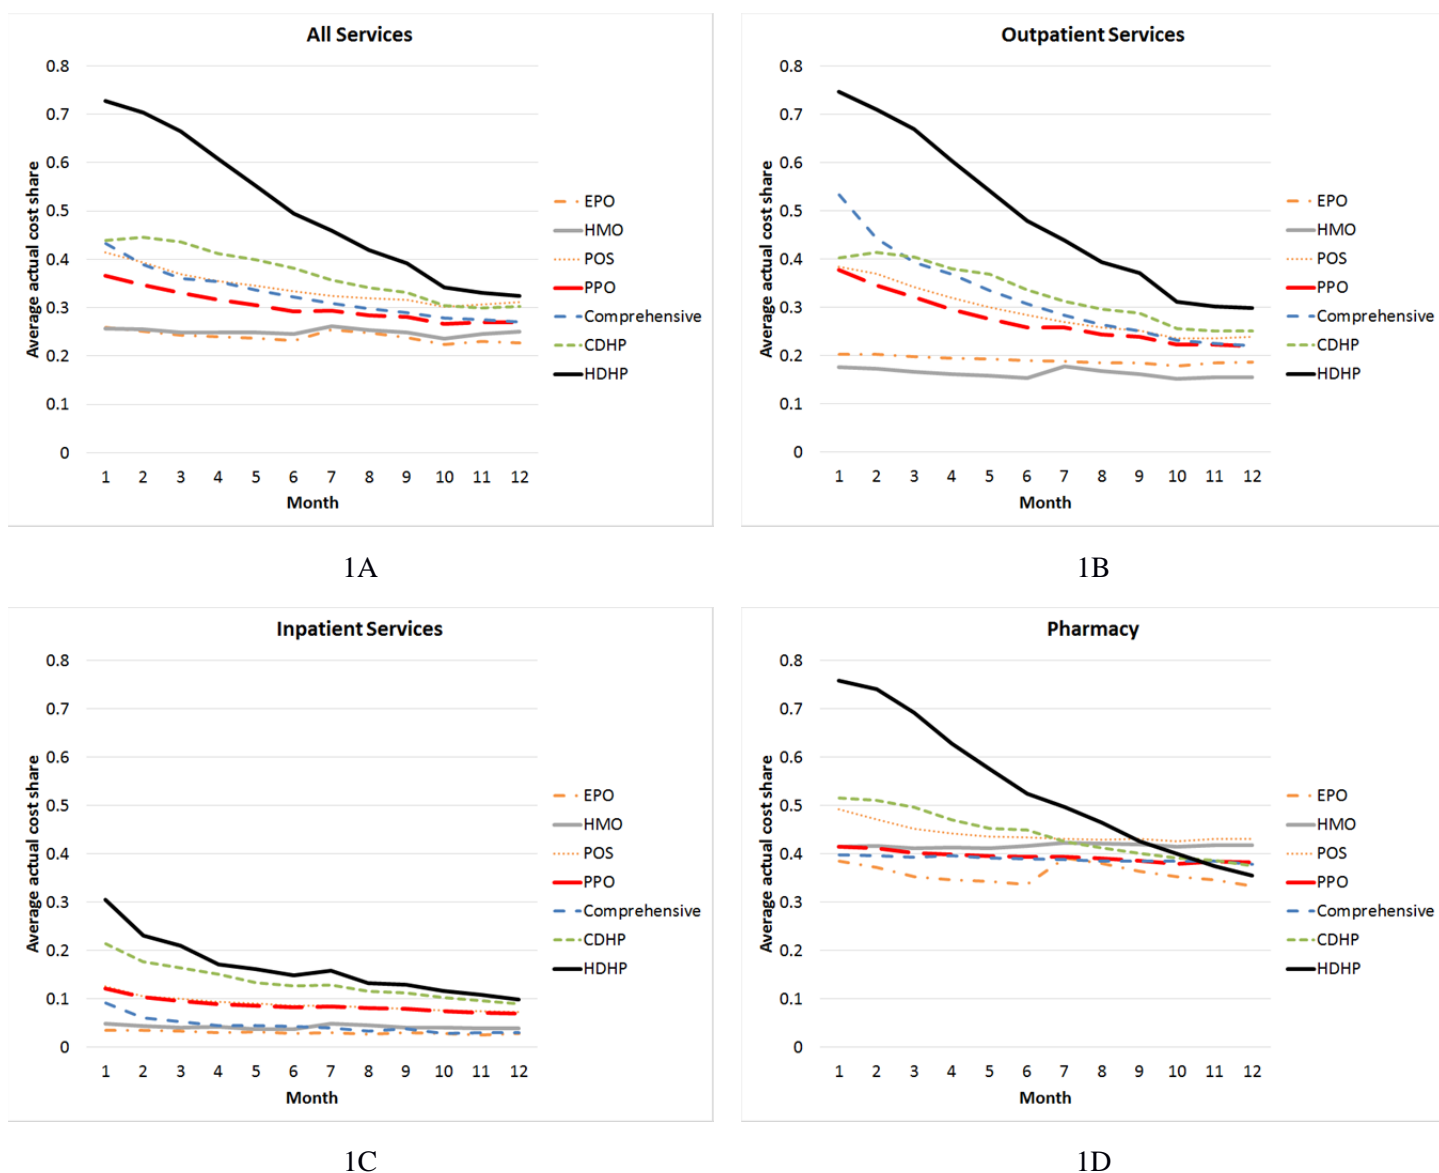


Fig 1. Average monthly cost share by plan type, 2008-2014

Note: Each panel plots the average actual cost share for the services indicated by month, averaged over the seven year period for each plan type for the 73 employer sample. People not using a service in the month are excluded.

we can see that outpatient spending shows the steepest decline, closely followed by pharmacy spending. Inpatient spending shows a relatively modest decline over months, and is particularly flat for HMOs and EPOs, as expected since they rely upon supply side controls more than demand side cost sharing to control costs. One surprise is that HMOs have a nearly constant and relatively high cost share for pharmacy spending (reflecting a nearly fixed copayment level) and charge a higher cost share than many other plan types by December.⁶

If health care spending is responsive to within-year declines in cost sharing, then it must be that spending increases during the year. And since cost sharing declines more in the HDHPs, CDHPs, and PPOs than other plan types, this increase should be greater for HDHPs and PPOs than in HMOs. We are not aware of any previous study that has documented this simple prediction. Corresponding to the average cost share by plan in Figure 1A to 1D, Figure 2A-2D presents average monthly spending in each of our seven plan types. Each diagram pools across all seven years and all the employers, and since we have a balanced sample with each person in all twelve months for each year in which they are in our analysis, simple means capture growth in spending over time.

The immediate pattern in all four plots is that costs are generally increasing in all plan types except for HMOs, where spending per month is nearly flat over time. The lowest line by December in every figure is the one for the HMOs, which is just below that of HDHPs. The line for CDHPs more or less follows the pattern of HDHP, low but strongly increasing. At the other extreme, costs are highest for the comprehensive health plan, which does not manage care, and also has the oldest enrollees and lowest average cost share. PPO spending is generally second highest, but also growing meaningfully over time.

The remaining panels of Figure 2 illustrate that mean costs are also noticeably increasing for outpatient spending and pharmaceutical spending. Spending growth is decidedly lower for inpatient spending, consistent with the slower changing monthly pattern for inpatient cost sharing, yet an upward trend in HDHP is still evident.

⁶ We speculate, but have not looked into the possibility, that HMOs do a better job at steering their enrollees toward generic drugs, so that even though the cost is lower, the share of drug costs is higher in HMOs than in plans that allow more use of branded drugs.

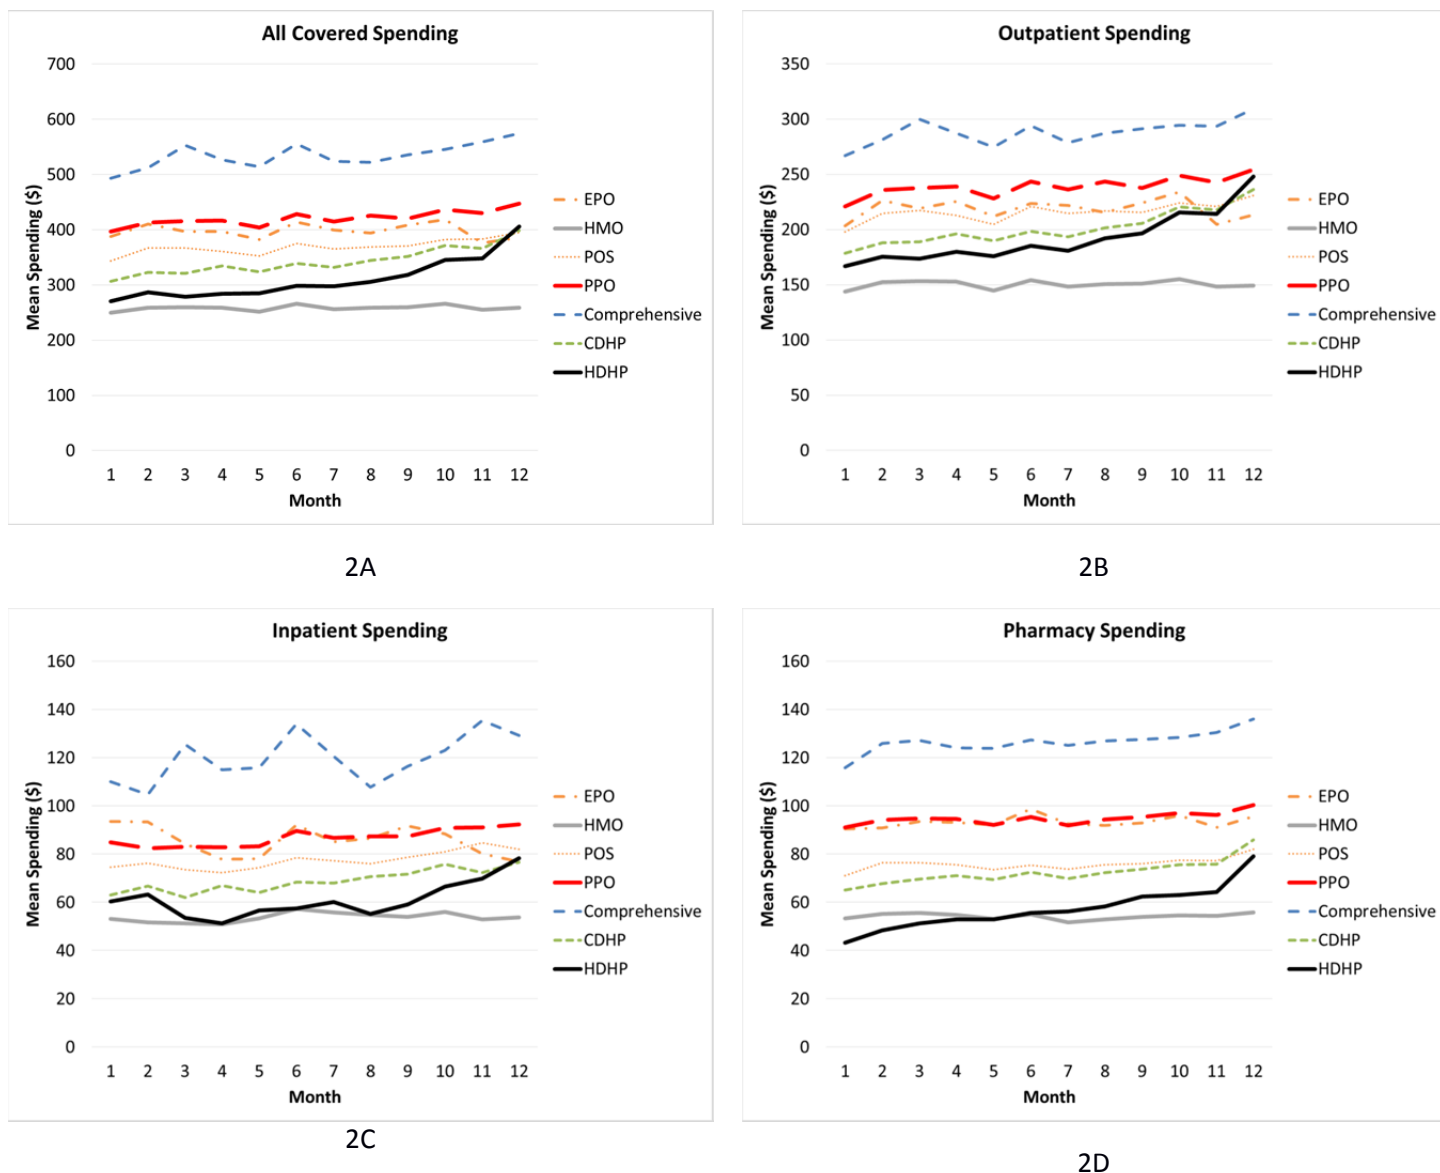


Fig 2. Average monthly spending by plan type, 2008-2014

Note: Each panel plots the monthly spending for the services indicated by month, averaged over the seven year period for each plan type for the 73 employer sample. People not using a service in the month are included as zero spending.

Figure 3 delves into these patterns more fully, using ten aggregated types of services for two plan types – HMO and PPO. Here it is easy to see that while cost shares are essentially flat across all service types in HMOs, they are sharply declining in HDHPs. Correspondingly, spending for many types of service are increasing sharply in HDHPs while essentially flat in HMOs. The two categories with the greatest increase in HDHPs over time are specialty care and

pharmaceuticals. After we introduce our empirical strategy for modeling price expectations, we use regression models to quantify the magnitudes of the changes over time, relate spending to cost share changes, and calculate demand elasticities.

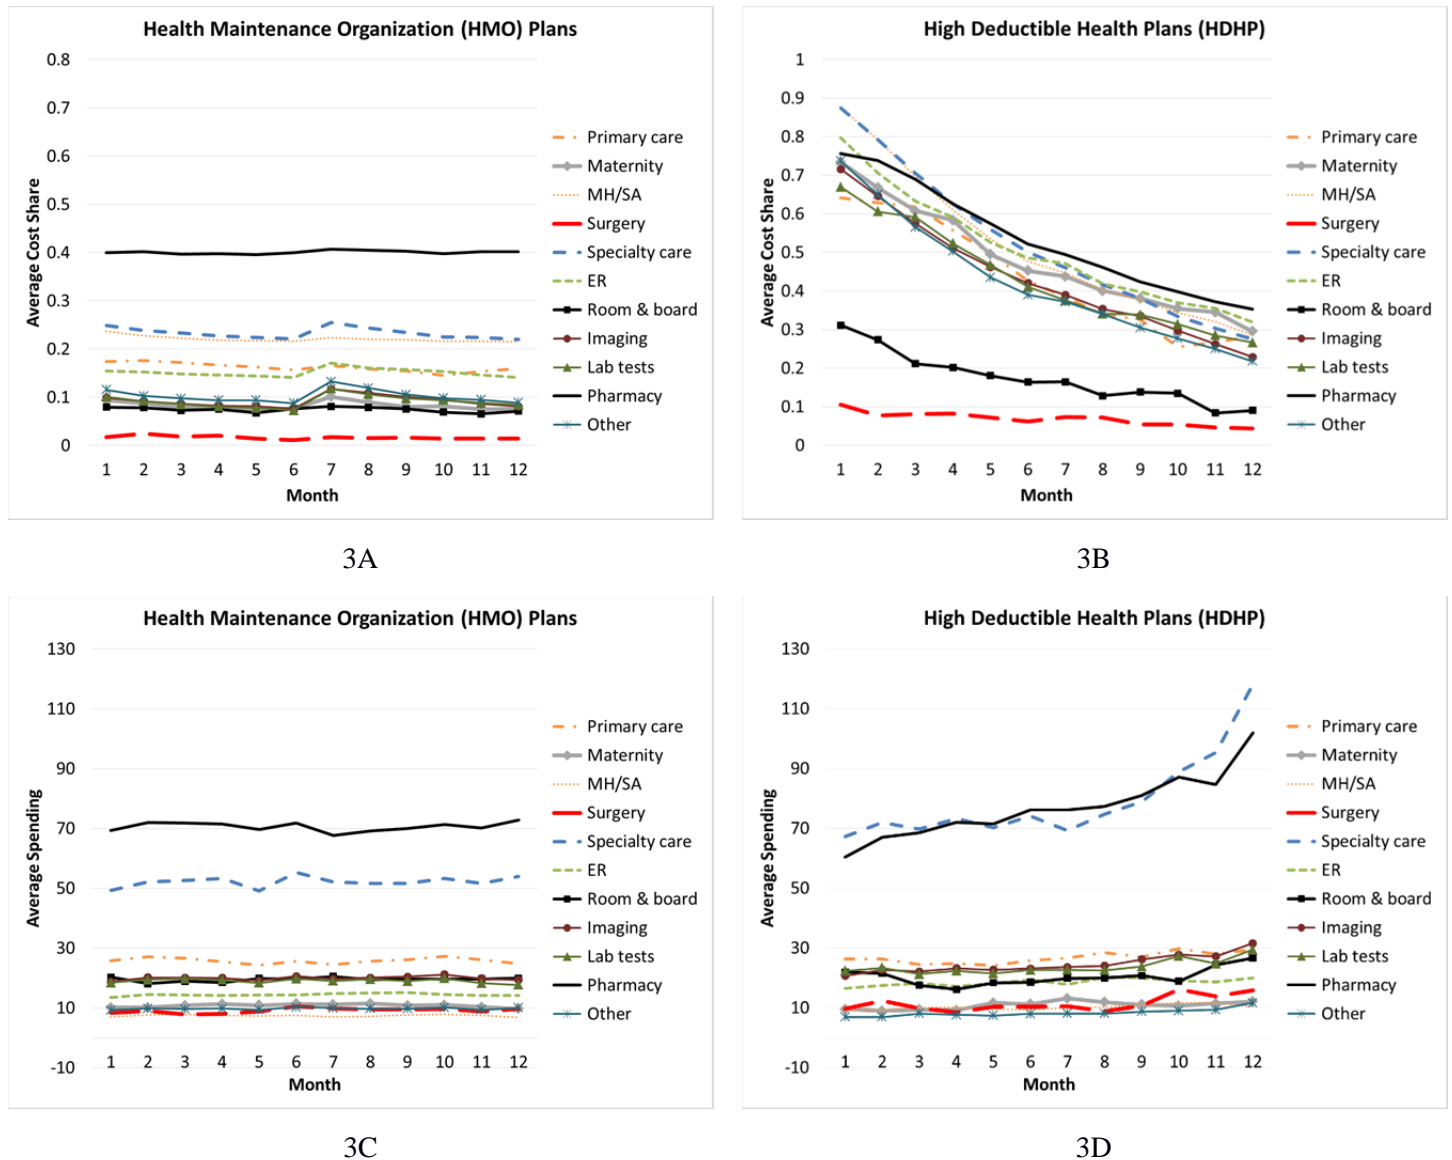


Fig. 3. Monthly average actual cost sharing and spending by aggregated Type of Service for HMOs and HDHP, 2008-2014

Notes: The top two panels (3A and 3B) plot the average actual cost shares by month for each of ten aggregated types of services, for HMOs and HDHP, respectively for the 73 employer sample. The bottom two panels (3C and 3D) plot the corresponding average spending per month in dollars for the same ten aggregated types of services.

3.3. *What prices (or cost shares) matter?*

For most goods, a consumer can choose how much to spend without worrying about how their own spending may change the marginal prices paid. Deductibles, coverage ceilings, and stoplosses introduce nonlinearities so that the rational consumer looks ahead when making purchasing decisions. In the original work by Keeler and colleagues (1988a, 1988b), the appropriate price to focus on is the expected end of year price. Ellis (1986) refined this by modeling how a risk-averse consumer should optimally decide using the end-of-year shadow price, which may be lower for a risk averse person with a declining block prices than the expected price. More recent studies, such as by Aron-Dine et al. (2012), Layton (2013), Einav et al. (2015), and Brot-Goldberg et al. (2015), as well as the numerous studies by behavioral economists in other settings, suggest that consumers are considerably more myopic than a rational consumer. Hence it is important to reflect some degree of myopia in our modeling.

In this paper, we consider two possible consumer prices (equivalently called cost shares) when modeling demand. Given that there is no cost share information for months in which a consumer did not have any spending, using the actual cost share is not possible without any further assumptions about consumer expectations. Furthermore, our data does not include general plan parameters, such as copay and deductibles that would allow us to predict what the cost share would be in the absence of spending. Some way of reflecting price expectations is needed to fill in all of the missing prices.

We consider two mechanisms of expectations, which likely provide upper and lower bounds on expected prices. The upper bound is a perfectly **myopic** expectations consumer, in which a consumer does not know the prices until after a visit is made and the cost share is actually observed. The actual monthly cost share is then used by the myopic consumer for decisions about each subsequent month up until another month with positive spending occurs, after which the myopic cost share is again revised. In this backward looking framework, it is easy to define prices after a visit is made, but an initial price is needed for each consumer before they make a visit, including those who never make a visit. Here we make the plausible assumption that consumers expect that the cost share for their first visit will be the average of the health plans cost share for all other consumers in their first month of seeking care. (If we knew the actual features of our health plans, we might assume that the initial price in a plan with a deductible was one, but our inference approach leads to much that outcome.) We overcome this problem by

assigning to individuals the average January cost share of those working for the same employer and under the same plan in that year.⁷

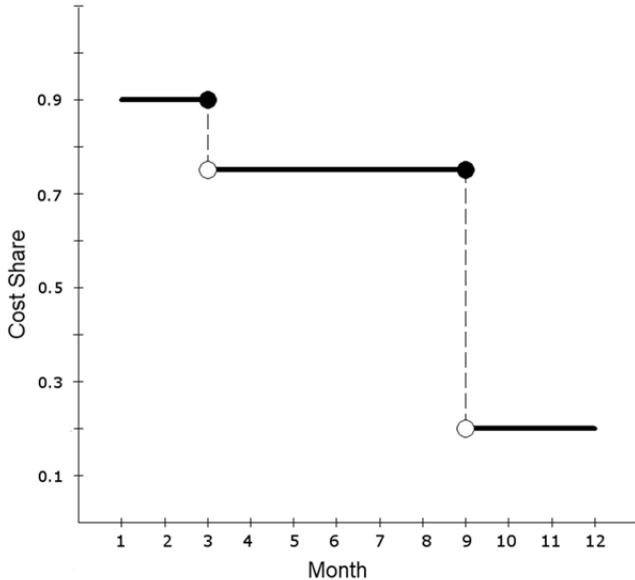
The other price expectations process we model is the optimal **forward-looking** consumer. The optimal forward-looking individual not only knows what the cost sharing will be at the time that services are chosen in a month, but also uses that price to make decisions prior to that first visit. If the year ends with a span of one or more months with no visits, we use the most recent cost share to complete the year. To close this model, we again need to model expectations for consumers who never make a visit. Specifically, we impute for all months the January average of those within the same employer and plan. Note that the extent to which the myopic and forward-looking expectation methods coincide depends on the amount of people who never make a visit in a year. The more months with no visits there are for a given type of service, the more similar the two imputation methods will be.

Figure 4 shows a schematic diagram of how expectations are assigned for one person who makes visits in March and September. In this figure, we depict a cost-share schedule for an individual with a deductible and a maximum out-of-pocket payment under a myopic view (panel (a)) and forward-looking (panel (b)). In panel (a), the myopic individual made no visits in January, and hence is assigned the average January cost share, which is 95% in this plan – almost the same as paying 100%, but lowered by the small number of people who exceed their deductible even in the first month. The myopic consumer uses this same cost share for February and March. In March, this individual obtains care and pays only 75% of the, so this is the new myopically expected cost share from April until and including the next visit month. Finally, in October, the consumer observes the 20% cost share from the previous month and uses this 20% cost share as the price of all subsequent months.

In panel (b), the forward-looking view, the consumer fully anticipates the March spending even in January, and uses that cost share for planning for January through March visits. This 75% rate could be because the consumer exceeded the deductible and the average cost share

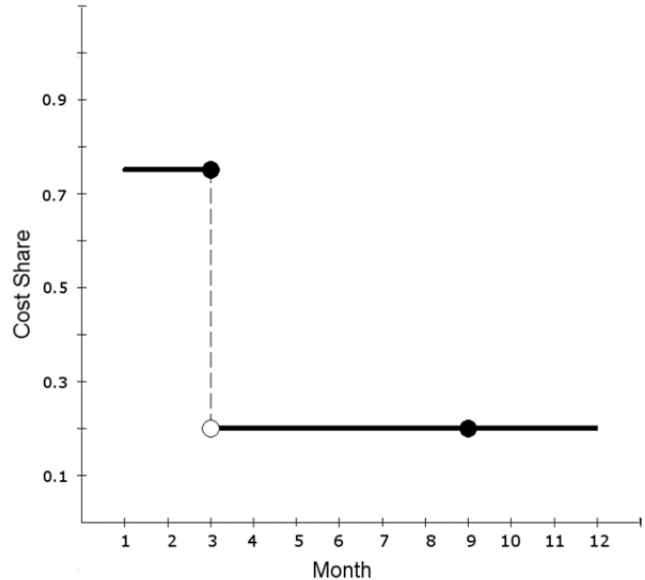
⁷ For the rare situation in which no one in a health plan make a visit for a given service in January, we drop all people in that plan for that year since we cannot impute their complete cost share expectations. Empirically, this affects only about 0.2% of the sample, mostly for rarely used services.

reflects some purchases at 100% and some at 20%. The 20% figure is the plausible rate for future visits. From April onwards, the individual expects the cost share to be the same on the margin as his March visit, at 20%, and this value remains in effect through September. Since no further spending occurs for the remainder of the year, it is also the effective forward-looking price through the end of the year.



Myopic uses the prior cost share.

4A



Forward-looking uses the next actual cost share.

4B

Fig. 4. Two models of expectations

Notes: Shown are the myopic and forward-looking expected cost shares for a single, hypothetical consumer who makes visits only in March and September, and who experiences an average cost share of .75 in March, and .2 in September. The average cost share for a person making their first visit in January for this employer*year*plan is assumed to be .90. The myopic consumer uses this price up through March, then lowers expectations after visits are made in March, and again in October. The forward-looking consumer anticipates the cost shares in March and September and uses them until new information comes along.

Figure 5 shows empirically the difference in the imputation methodology for three categories of services and three different plan types – HMO, PPO, and HDHP. HDHPs and HMOs are on opposite ends of the spectrum of the monthly cost share variation. Indeed, since the average cost sharing in HMOs is nearly constant from the beginning to the end of the year, the variation is nearly zero, making the imputations very close to the actual values. In sharp contrast, HDHPs have sharply declining actual, myopic and forward-looking prices. Both the myopic and forward-looking are above the actual values for most of the year. This reflects the mass of individuals that never make a visit and therefore always have a high price expectation.

This is particularly striking in the case of inpatient spending, where cost share imputations are nearly constant over time and shows very little variation across the year. The lack of monthly variation within a plan in inpatient cost sharing is a precursor of why we find it impossible to estimate the impact of cost share on inpatient spending and other similar services with precision even in our extremely large sample.

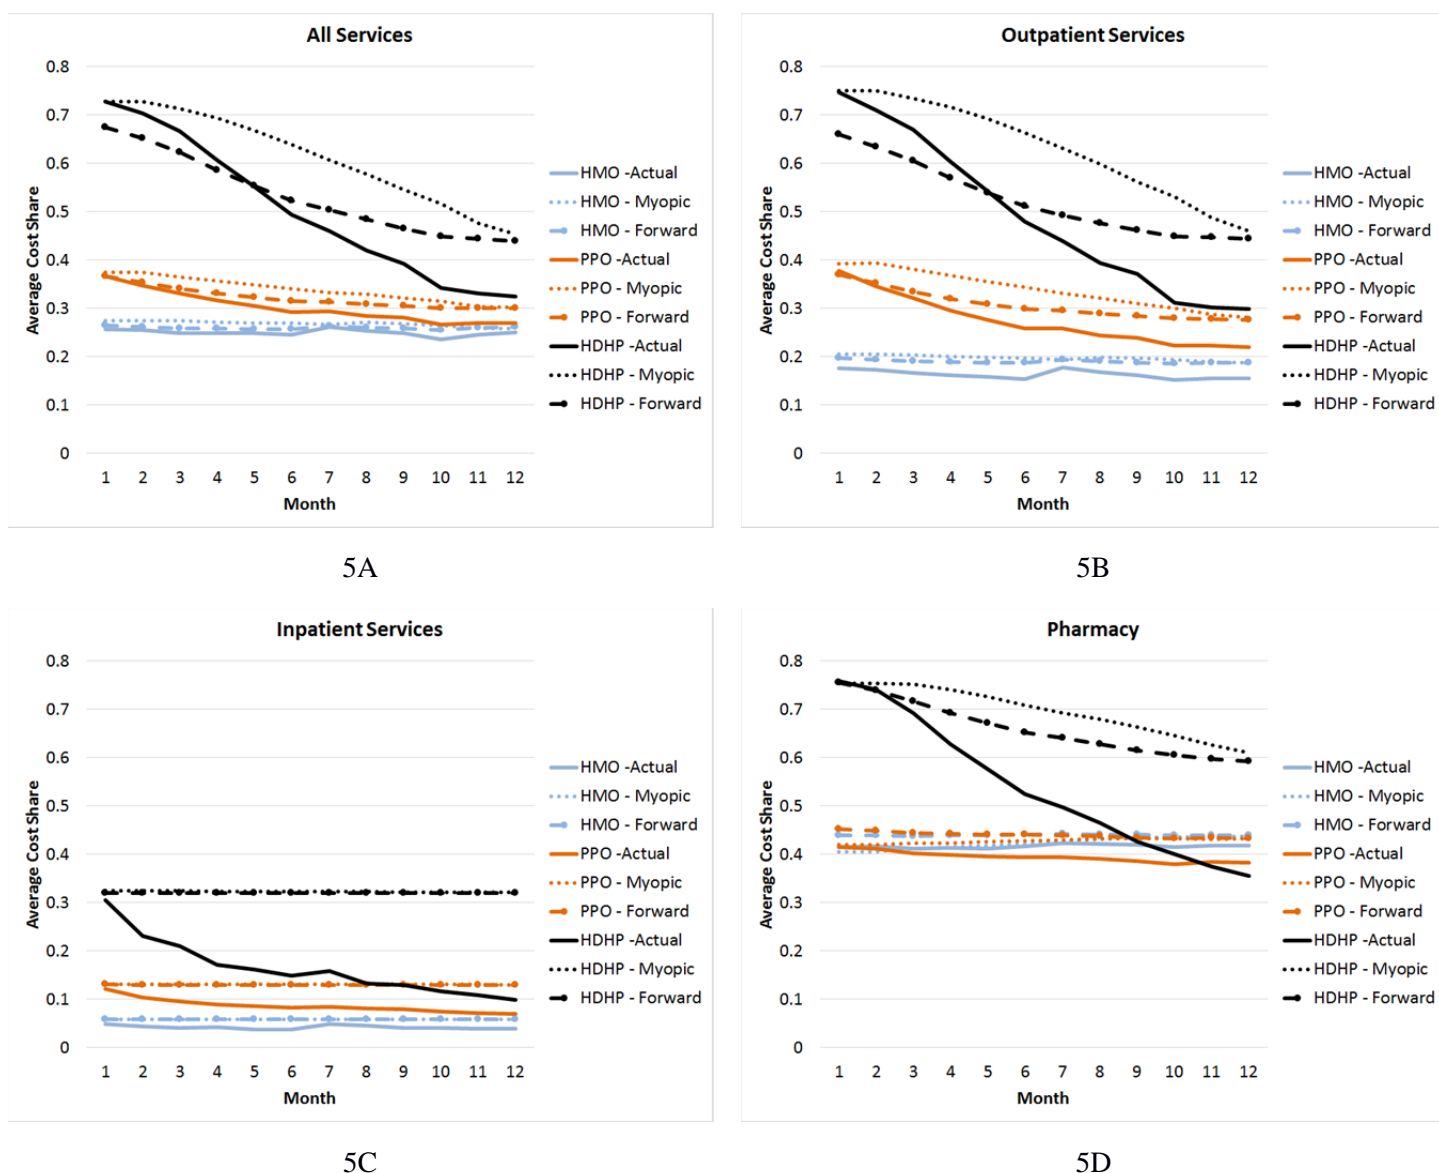


Fig. 5. Cost share imputations for HMOs, PPOs and HDHPs by Type of Service

Notes: Each panel shows the actual, myopic, and forward-looking cost shares for three plan types: HMOs, PPOs, and HDHPs. The four panels correspond to All, Outpatient, Inpatient and Pharmacy spending.

4. Empirical strategy

Since cost shares, including our myopic and forward-looking cost shares, reflect the actual experience of a consumer and hence is endogenous, using them in OLS for estimation will lead to biased estimates. (OLS regressions illustrating this bias are presented in appendix tables A-2 and A-3.) To be a valid instrument, we need the variable to be correlated with this individual price, but not correlated with any demand or supply side features that may change over time. Note that any fixed employer, market or individual characteristics will be absorbed by our fixed effects, so we are looking for instruments that change over the twelve month calendar year. Whereas the previous studies of Eichner (1998) and Kowalski (2016) used exogenous, individual-level shocks (injuries and other family member spending decisions), we use instead the monthly average cost shares for consumers in the same employer*year*family/single*plan as our instrument. To purge this measure of any contamination by a given consumer's experience, for each person we use the "all but one" average, which excludes that person's own contribution to the average. Hence if there are N people in a plan for a given month, we use the average of the N-1 other people in that plan. This measure is exogenous to the individual's own decisions, but summarizes well the average time path impact of deductibles and copayments across months within a given plan.

One of the things we worried about was not whether our IV was sufficiently correlated with the cost shares of interest (they are) but whether our actual cost shares are correlated with other spending in ways other than through demand effects. For example, it is possible that sicker people not only gravitate toward plans with higher cost sharing (they do) but so much so that their average cost shares were meaningfully higher in one plan than another. We therefore generated and examined Figures 6A-6D for all services and for our three broad categories of service. The horizontal axis plots intervals of prospective DxCG relative risk scores (RRS) for each person based on prior-year diagnostic information. So the highest value RRS correspond to being seven times the average spending, while the lowest RRS interval is .1, which is only 10 percent of the population average. Figure 6A shows that as expected the myopic and forward-looking prices have a sharp downward slope: sicker people are more likely to exceed any deductibles or stoplosses and pay lower cost shares. But the modestly downward slope for the employer mean actual cost share suggests that overall there is only a weak correlation between our instrument and the average health status of enrollees: The slight downward slope does

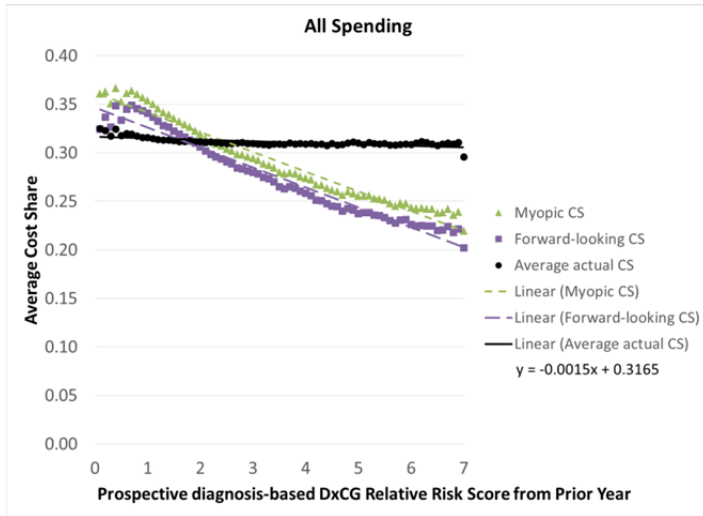
suggest that enrollees who are sicker tend to work at firms that on average have more generous coverage, so that their average cost share is slightly lower than the population average. We would have liked it if the black line for the employer means (our instruments) had zero slope rather than a negative slope with regard to this RRS, but this suggests our estimated demand elasticities will be slightly too high rather than too low. We interpret this analysis as saying that our IV strategy is acceptable, although not perfect.

5. Empirical results

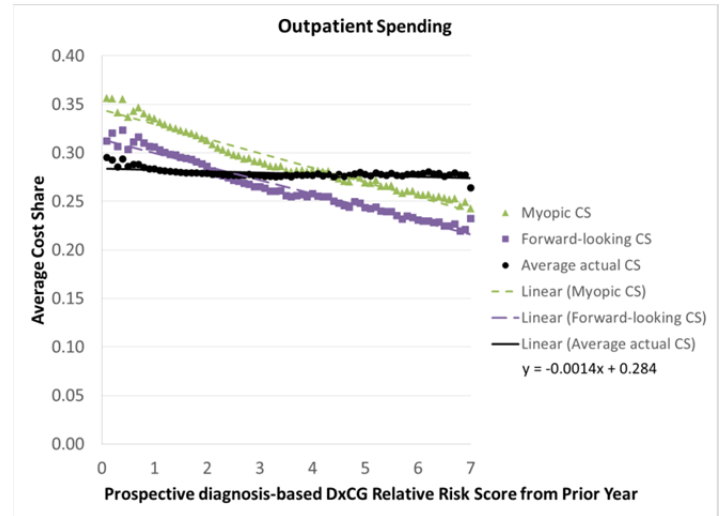
Our estimation approach is to estimate the total effect of cost sharing on service level spending using a linear regression model, as in Equation (2). Our preliminary analysis showed that linear models worked poorly since spending is highly skewed: even with our 172 million person months of data, we prefer the log linear specification, and as noted above, we used log of spending plus one as our dependent variable. (Aron-Dine et al. (2013) used the same specification.)

Table 4 presents the IV results from estimating model (2) on overall spending, on three broad categories of service (outpatient, inpatient, and prescription drugs), and on 11 finer types of service, using employer mean cost shares for a month as the instrument. Each row in the table is from a different TSLS regression that include also 12 monthly time dummies and absorbed individual*year fixed effects. Cluster corrections for all regression models were used to control for the fact that the employer*family*plan*year combination is much smaller than the number of individuals. Separate estimates were generated for myopic and forward-looking cost shares. Our overall demand elasticity is estimated to be -0.33 using the myopic cost share expectations and -0.31 for forward-looking cost share expectations. We view these as plausible and consistent with the previous literature.⁸

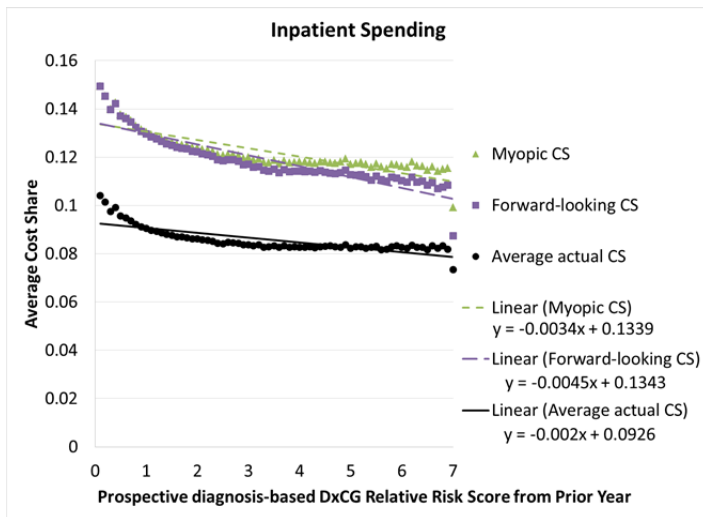
⁸ Kowalski (2016) used a different IV approach to estimate the overall demand elasticity of -1.49 at the median percentile.



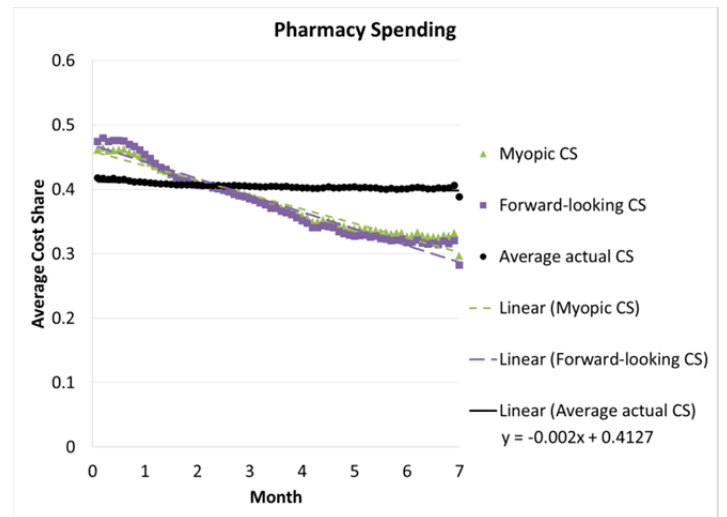
6A



6B



6C



6D

Fig. 6. Plot of intervals of risk scores versus monthly cost shares for three price measures, all spending, outpatient services, inpatient services, and pharmacy

Note: Each point in each diagram plots the average cost share for all person-month observations with a risk score in the 0.1 range of risk scores ranging from .1 to 7.0. Three different cost shares are plotted: the myopic and forward-looking cost shares, which are the individual level prices, and the average “all but one” employer cost share for a plan month, which is our instrument. The figures illustrate that while both myopic and forward prices decline sharply with higher risk scores, employer average costs in a month, which pool people with diverse risks, show very little decrease across risk scores. The modest declines imply there is some degree of bias upward in our demand elasticities, since sicker consumers (higher RRS) do have marginally lower average employer cost shares. N=172,122,188.

Table 4
IV regression results

<i>Model: log(Y+1)</i>	Myopic cost share			Forward-looking cost share		
	Coeff.	S.E.	Elasticity	Coeff.	S.E.	Elasticity
All	-0.95	(0.048)	-0.33	-0.95	(0.048)	-0.31
Outpatient	-0.61	(0.036)	-0.20	-0.64	(0.038)	-0.20
Inpatient	-1.74	(0.217)	-0.24	-13.51	(1.730)	-1.82
Drugs	-0.74	(0.038)	-0.33	-0.83	(0.042)	-0.38
Primary care	-0.55	(0.040)	-0.17	-0.72	(0.052)	-0.19
Maternity	-0.20	(0.056)	-0.06	-0.29	(0.082)	-0.09
MH/SA	-0.42	(0.036)	-0.18	-0.64	(0.055)	-0.27
Surgery	-2.29	(1.462)	-0.14	-70.29	(45.589)	-4.21
Specialty care	-0.58	(0.028)	-0.24	-0.74	(0.035)	-0.30
ER	-0.09	(0.032)	-0.03	-0.37	(0.136)	-0.12
Room and board	-1.15	(0.251)	-0.16	-6.65	(1.441)	-0.92
Imaging	-0.48	(0.025)	-0.15	-0.97	(0.051)	-0.28
Lab tests	-0.35	(0.026)	-0.11	-0.51	(0.038)	-0.15
Pharmacy	-0.75	(0.041)	-0.33	-0.85	(0.046)	-0.37
Other	-0.43	(0.040)	-0.16	-0.96	(0.090)	-0.35

Note: N = 172,122,188 individual months. Each row is a different regression using the log(spending plus 1) as the dependent variable. Each model includes individual*year fixed effects, 12 monthly time dummies and uses employer's average cost share for that year*plan*family*month, without the own individual, as an instrument. Cluster corrected standard errors use the employer*year*month*family*plan for clusters.

Our coarsest measure of type of service shows some heterogeneity between outpatient, inpatient and pharmacy claims. Our model estimates an outpatient elasticity of -0.2 for both the myopic cost share expectation and forward-looking. Pharmacy is even more elastic, at -0.33 (myopic) and -0.38 (forward-looking). Einav et al. (2016) estimate a mean drug elasticity of -0.24 for a Medicare Part D, a population older than our sample. ⁹Inpatient spending is estimated

⁹ They also estimate an overall drug elasticity of -0.037. However, this elasticity is defined as the probability of filing a claim for any drug, while ours is defined in terms of spending, which also reflects substitution effects between drugs. For this reason, we compare our estimates with the mean drug elasticity, which allows for substitution, rather than the overall value of -0.037.

to have a statistically significant elasticity of -0.24 using myopic versus an implausible -1.82 for the forward-looking prices. This is a symptom of almost all of estimates of elasticities for relatively rare, inpatient based services: elasticities are often implausibly large and generally imprecise. As previously pointed out, we believe this occurs because our IV approach fails to find a meaningful within-year variation in prices for these services.

Turning now to elasticities across the broad definition of types of services, we see that pharmacy and specialty care have slightly higher elasticities than other disaggregated services. Emergency room (ER) spending has an extremely low elasticity of -0.03 using myopic prices, as does maternity, consistent with our expectations. The forward-looking price elasticities are considerably larger once more detailed service types are used. Demand response for mental health and substance abuse (MH/SA) services are of intermediate elasticities (-0.18 (myopic) and -0.27 (forward-looking)). This may reflect the current trend that most people getting professional mental health treatment are getting drugs only or receiving MH/SA treatment in primary care (McGuire 2016), which is consistent with the low level of spending on specialty mental health care observed in our data. Inpatient surgery is not statistically different from zero.

Table 5 presents the results of the same model for 33 detailed types of services. We intentionally included both relatively common groups, such as non-specialty visits and pharmacy, along with relatively rare groups such as six specific types of imaging separately. With our large sample, we find that 26 of the 33 detailed services are statistically significantly different from zero at the 5% level. Again, for simplicity, we discuss only the results for myopic expectations. At a detailed level, among our statistically significant results, we find larger than average demand responsiveness for specialty visits (-0.22), MRIs (-0.21) and pharmacy (-0.33). Relatively inelastic (with elasticities below -0.1 using myopic expectations) are prevention, maternity, non-surgery inpatient procedures, emergency room, mammograms, ultrasounds, and specialty drugs.

Table 5

IV regression Results, by detailed Type of Service (TOS)

<i>Model: log(Y+1)</i>	Myopic cost share			Forward-looking cost share		
	Coeff.	S.E.	Elasticity	Coeff.	S.E.	Elasticity
Non-specialty visits	-0.45	(0.036)	-0.16	-0.56	(0.044)	-0.17
Home visits	-0.02	(0.189)	0.00	-0.31	(0.899)	-0.07
Prevention	-0.22	(0.078)	-0.01	-3.19	(1.109)	-0.19
Maternity	-0.20	(0.056)	-0.06	-0.29	(0.082)	-0.09
Mental health	-0.42	(0.036)	-0.18	-0.64	(0.056)	-0.27
Substance abuse	-0.14	(0.118)	-0.05	-0.34	(0.315)	-0.13
Surgical procedures	1.05	(1.363)	0.06	32.11	(30.701)	1.73
Surgical supplies/devices	-2.53	(0.964)	-0.17	-61.75	(25.179)	-4.08
Nonsurgery IP procs	-0.38	(0.151)	-0.06	-3.42	(1.271)	-0.54
Specialty visits	-0.56	(0.027)	-0.22	-0.77	(0.038)	-0.29
Dialysis	0.11	(0.549)	0.01	0.08	(1.018)	0.01
PT, OT, speech therapy	-0.22	(0.030)	-0.10	-0.41	(0.055)	-0.18
Chiropractic	-0.23	(0.031)	-0.15	-0.30	(0.040)	-0.19
Hospice	46.91	(23.262)	3.92	-7.08	(54.025)	-0.59
Emergency room	-0.09	(0.032)	-0.03	-0.37	(0.136)	-0.12
Room board - Surgical	-1.09	(0.482)	-0.13	-16.58	(7.204)	-1.95
Room board - Medical, other	-0.44	(0.279)	-0.07	-2.83	(1.653)	-0.43
CAT scan	-0.34	(0.036)	-0.10	-1.84	(0.196)	-0.56
Mammograms	-0.70	(0.094)	-0.07	-13.81	(1.839)	-1.36
MRIS	-0.69	(0.037)	-0.21	-5.34	(0.288)	-1.61
PET scans	-0.85	(0.350)	-0.13	-4.69	(2.295)	-0.71
Radiology - diagnostic	-0.28	(0.017)	-0.10	-0.72	(0.044)	-0.25
Radiology - therapeutic	0.36	(0.582)	0.05	-0.50	(2.902)	-0.07
Ultrasounds	-0.26	(0.019)	-0.09	-1.05	(0.076)	-0.38
Diagnostic services	-0.35	(0.025)	-0.10	-1.04	(0.073)	-0.28
Laboratory other	-0.33	(0.026)	-0.11	-0.50	(0.040)	-0.15
Pharmacy	-0.74	(0.038)	-0.33	-0.83	(0.042)	-0.38
Facility-based pharmacy	-0.41	(0.094)	-0.10	-1.84	(0.425)	-0.45
Specialty drugs, injections	-0.12	(0.024)	-0.05	-0.26	(0.051)	-0.09
Non-surg supplies/devices	-0.44	(0.027)	-0.18	-1.34	(0.082)	-0.54
DME	-0.26	(0.028)	-0.12	-0.68	(0.074)	-0.32
Transportation	-0.22	(0.170)	-0.05	-3.04	(2.668)	-0.64
Other	-0.42	(0.056)	-0.15	-1.09	(0.144)	-0.37

Note: N = 172,122,188 individual months. Each row is a different regression using the log(spending plus 1) as the dependent variable. Each model includes individual*year fixed effects, 12 monthly time dummies and uses employer's average cost share for that year*plan*family*month, without the own individual, as an instrument. Cluster corrected standard errors use the employer*year*month*family*plan for clusters.

Although it may seem that 172 million observations is very large, given the infrequent usage in monthly data, there is remarkably little power to detect patterns for cost shares of rarely used services, particularly when they are relatively expensive. Table 3, which provides summary statistics by detailed types of services, identifies the services with very low rates of positive spending but high conditional spending: home visits, dialysis, and hospice are examples of this. Most people with spending on these types of service are likely to be associated with high annual spending, so that many people will exceed any reasonable deductible or stoploss, and hence face very little price variation. This is also a good argument for why cost sharing on such services will be an ineffective cost containment tool, since it imposes financial risk with little effect on spending on these services.

5.1. Extensions of the basic model

In this section, we re-estimate the model for the total spending and for various partitions of our sample to consider whether demand response is higher or lower for certain identifiable groups. We also compute arc-elasticities instead of the point elasticities estimated so far. We then conduct a number of sensitivity analyses.

5.2. Demand elasticities by selected subsamples

By gender and age. In this section we examine price elasticities of demand of total spending and the three broad groups of services – outpatient, inpatient, and pharmacy – for various subgroups of our sample. Table 6 presents results for five different partitions of individual-years in our estimation sample using the myopic expectations formulation. First, there is no appreciable difference in elasticities between males and females overall in our sample. Second, across all three service types, spending on children is significantly more inelastic than spending on older individuals, with the elasticities growing monotonically with age for all categories except for inpatient spending, which is imprecisely estimated for age groups 0-5 and 6-20¹⁰.

¹⁰ The coefficient estimates of inpatient spending are not statistically significant on these two age groups.

Table 6

Elasticities results by demographic and plan subgroups

Model: $\log(Y+1)$, myopic expectations, IV results					
	% of sample	All spending	Outpatient	Inpatient	Pharmacy
Gender					
Male	48.9%	-0.34	-0.20	-0.19	-0.33
Female	51.1%	-0.32	-0.20	-0.26	-0.33
Age group					
0 to 5	6.3%	-0.12	-0.05	-0.21	-0.22
6 to 20	24.2%	-0.21	-0.09	-0.14	-0.30
21 to 45	34.8%	-0.37	-0.22	-0.31	-0.31
46 to 64	34.7%	-0.40	-0.28	-0.20	-0.36
Risk score					
0.0 to 0.99	69.3%	-0.32	-0.18	-0.02	-0.28
1.0 to 1.99	18.2%	-0.36	-0.25	-0.06	-0.34
2.0 to 3.99	9.3%	-0.29	-0.22	-0.24	-0.35
4 or more	3.2%	-0.21	-0.15	-0.13	-0.31
Plan type					
HMO	12.7%	-0.30	-0.08	0.03	-0.43
PPO	63.0%	-0.29	-0.18	-0.27	-0.25
HDHP	1.8%	-0.62	-0.59	-0.88	-0.58
Time period					
2008-2009	31.3%	-0.20	-0.12	-0.18	-0.10
2010-2012	49.5%	-0.38	-0.26	-0.21	-0.34
2013-2014	19.2%	-0.29	-0.18	-0.28	-0.42
Plan coverage					
Single	22.1%	-0.43	-0.28	-0.10	-0.39
Family, employee	24.7%	-0.42	-0.27	-0.28	-0.35
Family, spouse	19.6%	-0.35	-0.23	-0.31	-0.32
Family, child	33.6%	-0.19	-0.09	-0.20	-0.28

Note: Table shows the estimated elasticities for all services and for three broad types of services by gender, age group, risk score ranges, selected plan types, time interval and plan coverage. Each elasticity is from a different IV regression. Each IV regression includes individual*year fixed effects and 84 monthly time dummies and uses employer's average cost share for that service in that year*plan*family*month as an instrument for the myopic cost share. Dependent variable is log of spending on that service plus one. Cluster-corrected standard errors use the employer*year*month*family*plan for clusters.

By risk scores. The third dimension evaluated is by intervals of risk scores, which suggests that the overall demand elasticity is higher for low risk score individuals relative to high risk scores. There is also a striking pattern across health plan type, in which HMO enrollees are significantly less price responsive than HDHP. This finding has implications for the work of Brot-Goldberg et al. (2015) which focused on changes in demand from a firm that significantly raised its deductibles, and found relatively large demand elasticities. This result is also consistent with the results of Einav et al. (2015) who document heterogeneity in Medicare pharmacy demand.

By time. The next dimension examined is changes over time, for which our estimates suggest the demand for pharmacy and inpatient are becoming more elastic between 2008 and 2014, while the demand for outpatient and total spending increased from 2008-09 to 2010-12 before declining slightly going into 2013-14. The final set of results in Table 6 relates to single versus family coverage. We find that single enrollees have elasticities of demand that are essentially the same as employees at family coverage, but both groups are more elastic than either spouses or children.

By employment characteristics. Table 7 continues the analysis of subgroups of our sample, using dimensions that relate to employment. Classifying employers by their one digit industry code, we see that enrollees working at a firm in transportation, communications, or utilities have the least elastic demand for overall services (driven by an inelastic demand for outpatient spending), followed by service industry both of which have less elastic demand than manufacturing (durable goods) or finance, insurance and real estate. Service industry has the lowest elasticity of inpatient spending, while spending on pharmacy is the most inelastic among people working at manufacturing (durable goods) industry. Furthermore, categorized by the type of salaries paid, the lowest elasticity workers are in unionized jobs, (whether salaried or hourly). Among nonunion workers, salaried workers are more elastic than hourly workers, except for pharmacy spending.

By high versus low cost share employers. The last set of results in Table 7 ordered employer-years by quartile according to their average overall cost share, so the lowest quartile firms are those with very little consumer cost sharing, while the highest quartile firms have relatively high

consumer cost sharing. Our analysis reveals that higher cost sharing firms also have more elastic enrollees, consistent with our finding that individuals in HDHPs are more responsive to cost sharing than individual not in such plan. This finding is consistent with a theory that consumers only start responding to cost sharing when such out of pocket spending is meaningful.

Table 7
Elasticities results by employment-related subgroups

Model: log(Y+1), myopic expectations, IV results					
	% of sample	All spending	Outpatient	Inpatient	Pharmacy
Industry					
Services	11.5%	-0.26	-0.13	-0.12	-0.51
Manufacturing, Durable Goods	14.3%	-0.40	-0.33	-0.19	-0.26
Finance, Insurance, Real Estate	12.3%	-0.39	-0.24	-0.24	-0.43
Transportation, Communications, Utilities	13.0%	-0.10	-0.03	-0.49	-0.41
Salary class					
Union (salaried or hourly)	14.0%	-0.27	-0.14	-0.11	-0.31
Salaried, non-union	26.3%	-0.33	-0.21	-0.30	-0.28
Hourly, non-union	19.4%	-0.28	-0.15	-0.26	-0.36
Other	40.4%	-0.35	-0.23	-0.19	-0.40
Employers by level of average cost sharing					
Lowest quartile	22.6%	-0.16	-0.08	-0.09	-0.24
Second quartile	34.3%	-0.27	-0.16	-0.22	-0.24
Third quartile	25.9%	-0.43	-0.24	-0.29	-0.33
Highest quartile	17.2%	-0.54	-0.67	-0.32	-0.50

Note: Table shows the estimated elasticities for all services and for three broad types of services by industry, salary class, and employer quartiles of costs share. For the last group, individual-years were divided into four quartiles according to a ranking of the average cost share of their employer in that year, such that the lowest quartile are at employers offering the lowest average cost share. Each elasticity in this table comes from a different IV regression. Each IV regression includes individual*year fixed effects, 84 monthly time dummies and uses employer's average cost share for that service in that year*plan*family*month as an instrument. Dependent variable is log of spending on that service plus one. Cluster-corrected standard errors use the employer*year*month*family*plan for clusters.

5.3. Arc-elasticities

So far, our analysis has focused on point elasticities based on the mean cost share for each type of service. Point elasticities may not give us a complete picture of demand responsiveness especially for the ranges of cost shares in which we do not observe services being used. Hence we also compute arc-elasticities for overall spending, three broad categories of service, and 11 finer types of services using the same cost share brackets as Aron-Dine et al. (2013), reported in Table 8.¹¹

The results show that elasticities vary substantially across cost share ranges, and using myopic and forward-looking cost shares gives us qualitatively similar results. Keeping everything else fixed, elasticities are generally larger for bigger cost share brackets. Overall, inpatient spending is the most elastic of the three types, driven mainly by surgery. This result is not captured through the point elasticities, as these types of services usually have low cost sharing (spending is usually high enough to put consumers past deductibles or even stoplosses), and thus we do not observe choices being made for these services on a wide range of cost share. Emergency room and maternity show up as the most inelastic services.

¹¹ Arc-elasticities are computed as $(S' - S) / ((S' + S) / 2)$, where S and S' are the exponential of the cost share multiplied by the regression coefficient. Note that we omit the adjustment of adding one to spending, otherwise arc-elasticities would have to been evaluated at some level of patient characteristics, which is not possible due to the individual*year fixed effects.

Table 8
Arc-elasticities

Arc elasticities	Myopic						Forward-looking					
	Free care to...			0.25 cost share to...		0.5 cost share to...	Free care to...			0.25 cost share to...		0.5 cost share to...
	0.25	0.5	0.95	0.5	0.95	0.95	0.25	0.5	0.95	0.5	0.95	0.95
Cost share	0.25	0.5	0.95	0.5	0.95	0.95	0.25	0.5	0.95	0.5	0.95	0.95
All spending	-0.12	-0.23	-0.42	-0.35	-0.55	-0.68	-0.12	-0.23	-0.42	-0.35	-0.55	-0.67
Outpatient	-0.08	-0.15	-0.28	-0.23	-0.36	-0.44	-0.08	-0.16	-0.30	-0.24	-0.38	-0.46
Inpatient	-0.21	-0.41	-0.68	-0.64	-0.93	-1.20	-0.93	-1.00	-1.00	-2.80	-1.71	-3.21
Pharmacy	-0.09	-0.18	-0.34	-0.28	-0.43	-0.53	-0.10	-0.20	-0.37	-0.31	-0.48	-0.59
Primary care	-0.07	-0.14	-0.25	-0.20	-0.32	-0.39	-0.09	-0.18	-0.33	-0.27	-0.42	-0.52
Maternity	-0.02	-0.05	-0.09	-0.07	-0.12	-0.14	-0.04	-0.07	-0.14	-0.11	-0.17	-0.21
MH/SA	-0.05	-0.11	-0.20	-0.16	-0.25	-0.31	-0.08	-0.16	-0.30	-0.24	-0.38	-0.46
Surgery	-0.28	-0.52	-0.80	-0.84	-1.14	-1.53	-1.00	-1.00	-1.00	-3.00	-1.71	-3.22
Specialty care	-0.07	-0.14	-0.27	-0.22	-0.34	-0.42	-0.09	-0.18	-0.34	-0.28	-0.44	-0.53
ER	-0.01	-0.02	-0.04	-0.03	-0.05	-0.06	-0.05	-0.09	-0.17	-0.14	-0.22	-0.27
Room and board	-0.14	-0.28	-0.50	-0.43	-0.65	-0.81	-0.68	-0.93	-1.00	-2.04	-1.68	-2.91
Imaging	-0.06	-0.12	-0.22	-0.18	-0.28	-0.34	-0.12	-0.24	-0.43	-0.36	-0.56	-0.69
Lab tests	-0.04	-0.09	-0.17	-0.13	-0.21	-0.26	-0.06	-0.13	-0.24	-0.19	-0.30	-0.37
Pharmacy	-0.09	-0.19	-0.34	-0.28	-0.44	-0.54	-0.11	-0.21	-0.38	-0.32	-0.49	-0.61
Other	-0.05	-0.11	-0.20	-0.16	-0.25	-0.31	-0.12	-0.24	-0.43	-0.36	-0.56	-0.69

5.4. Sensitivity analysis

Table 9 presents results from including multiple prices so as to estimate cross price elasticities along with own price elasticities, and also conducts a number of specification tests on our framework. Since they are more stable and robust, we focus here on models using myopic price expectations. The first block of results repeats our benchmark results which are for myopic price expectations using the plan*single/family*year*month mean cost share as an instrumental variable. The next three blocks (2-4) estimate not only the own price effect but also cross price effects for three broad sets of services. The estimates suggest that the own price effects when other prices are omitted, are not highly dissimilar to those when cross price terms are included. Detailed estimates on the cross price effects (presented in appendix Table A-4) find that the effects of outpatient cost share on inpatient spending and on pharmacy spending are not statistically significant. As before, estimates for inpatient services are unstable and the point

estimates are relatively large in many specifications. We do not have any explanation for why they appear to be statistically significant here.

Table 9

Sensitivity analysis of alternative specifications

Benchmark	All spending	Outpatient	Inpatient	Pharmacy
	-0.33	-0.20	-0.24	-0.33
Cross-price elasticities		Outpatient	Inpatient	Pharmacy
Outpatient cost share		-0.12	-0.0001	0.01
Inpatient cost share		-3.33	-0.26	-2.66
Pharmacy cost share		-0.17	0.0032	-0.28
Alternative models	All spending	Outpatient	Inpatient	Pharmacy
January and December dropped	-0.32	-0.20	-0.33	-0.28
Alternative IV: total cost share	-	-0.26	-0.22	-0.38
Log(Y+0.1)	-0.43	-0.28	-0.29	-0.43
Log(Y+10)	-0.22	-0.13	-0.18	-0.21

Note: Table shows the robustness of estimated elasticities for all services and for three broad types of services to including cross-price effects, dropping the first and last months of observations, using employer's average cost share for all services as an instrument to cost share for each service, and to alternative log transformations of the dependent variable. Each IV regression includes individual*year fixed effects, 84 monthly time dummies and uses employer's average cost share in that year*plan*family*month as an instrument. Dependent variable is log of spending on that service plus one. Cluster-corrected standard errors use the employer*year*month*family*plan for clusters.

Table A-4 further examines various alternative specifications to our main results. There we show that our results are robust to dropping January and December from our estimation and using only ten months of data for each person year. This is relevant in that there is some hint in the graphical results that there may be anticipation effects in December, with higher spending at the end of the year and somewhat lower spending in January. Spending in December could also be motivated by medical spending accounts which can have a “use it or lose it” effect. (The same should not be true in the health savings accounts where balances can be rolled over with no

penalty.) Our results find no meaningful change in any of the demand elasticities from excluding these terminal months.

Table A-4 also considers using a different instrument for the broad sets of spending, in that we use the employer's all service monthly average cost share as the IV rather than the average for that particular service. For our approach this should still be a valid if somewhat weaker IV, since it is more weakly correlated with the right hand side. Here the coefficient and elasticity for outpatient spending becomes smaller and is statistically significantly different, while there are no statistically significant changes in the inpatient and pharmacy results.

The final segment of Table A-4 examines the implications of using $\log(Y+0.1)$ or $\log(Y+10)$ instead of $\log(Y+1)$. Regardless of the specification, outpatient spending is the least responsive to prices, followed by inpatient and pharmacy spending. What is less satisfying is that the estimated demand elasticities are moderately sensitive to the constant that is added on, with all elasticities roughly doubling in magnitude as the additive constant changes from 10 to 0.1. This is a known problem with using the $\log(Y+1)$ specification, although the sensitivity results for changing this constant are not always reported in other empirical studies using this approach.

6. Conclusions

This paper develops a new IV approach for estimating demand responsiveness using big data and highly disaggregated types of service. We document and take advantage of the considerable within-year variation in cost shares in many health plan types, which when combined with plans that have flat cost sharing creates a nice setting for estimation of demand elasticities. It seems not to be widely recognized that downward cost sharing trends during the calendar year must imply upward spending trends on services if such use is not fully anticipated, or consumers are myopic. We show that these patterns are strong across many types of service, showing that within year price responsiveness is significant.

Our IV strategy leads to elasticity estimates that are plausible and consistent with other estimates from the literature for services where a single month of use does not typically put a person over normal deductible levels. Our IV approach works less well for expensive procedures like hospice care, inpatient surgery, room and board spending, and maternity, where the time of the year matters less since the consumer will invariably exceed almost any deductible. While not perfect, our IV has the advantage of being widely available and easy to use. OLS estimates, in

contrast are often in the wrong sign, and our results suggest that studies using only a single employer, or using only people who are in high deductible plans can be seriously biased, since they are not representative of the full population of privately insured in our sample.

Although we estimate models using two different forms of expectations: myopic expectations where consumers do not anticipate future spending, and forward-looking expectations, where they do, we find relatively modest differences in our elasticities once we ignore statistically insignificant ones. One hypothesis for this is that our two measures are highly correlated when consumers use a service multiple times per year, and since we use an IV strategy, it corrects for expectation errors.

Our approach holds great promise in potentially providing a new instrument – employer average monthly cost share on a service of interest – which can be used in other studies looking for an instrument for rates of utilization of a service. For example, consider procedure ABC, lab test DEF, or drug GHI that are sometimes used only for a patient with condition 123. If consumers are diagnosed with condition 123 at different times in the year, and face changing cost shares over time, then our approach could be used as a reasonable instrument for assessing the effectiveness of each of these procedures, tests, or drugs for treating condition 123. Given how valuable good instruments are, this approach could be as useful as the estimates of demand response that we have generated so far.

References

- Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein. 2013. "The RAND Health Insurance Experiment, Three Decades Later." *Journal of Economic Perspectives*, 27(1): 197-222.
- Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark R. Cullen. 2012. "Moral Hazard in Health Insurance: How Important is Forward Looking Behavior?" NBER Working Paper 17802.
- Ash, Arlene S., Randall P. Ellis, Gregory C. Pope, et al. 2000. "Using Diagnoses to Describe Populations and Predict Costs." *Health Care Financing Review*, 21(3): 7-28.
- Baker, Laurence, Kate Bundorf, Anne Royalty, et al. 2007. "Consumer-Oriented Strategies for Improving Health Benefit Design: An Overview." Rockville (MD): Agency for Healthcare Research and Quality (US); (Technical Reviews, No. 15.) Available from: <http://www.ncbi.nlm.nih.gov/books/NBK44061/>
- Brachet, Tanguy. 2007. "Computing Clustered Standard Errors for Two-Stage Least Squares in SAS." (software) <http://works.bepress.com/tbrachet/2/>.
- Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. November 2, 2015. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." Working Paper.
- Cameron, Colin and Douglas L. Miller. 2015. "A Practitioner's Guide to Cluster-Robust Inference." *Journal of Human Resources*, 50(2): 317-372.
- Duarte, Fabian. 2012. "Price Elasticity of Expenditure across Health Care Services." *Journal of Health Economics*, 31(6): 824-841.
- Eichner, Matthew J. 1998. "The Demand for Medical Care: What People Pay Does Matter." *American Economic Review*, 88(2): 117-121.
- 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 Paul Schrimpf. 2015. "The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D." *The Quarterly Journal of Economics*. doi: 10.1093/qje/qjv005
- Einav, Liran, Amy Finkelstein, and Maria Polyakova. 2016. "Private Provision of Social Insurance: Drug-Specific Price Elasticities and Cost Sharing in Medicare Part D." NBER Working Paper 22277
- Ellis, Randall P. 1986. "Rational Behavior in the Presence of Coverage Ceilings and Deductibles." *The Rand Journal of Economics*, 17(2): 158-175.
- Ellis, Randall P., and Thomas G. McGuire. 1986. "Cost-Sharing and Patterns of Mental Health Care Utilization." *Journal of Human Resources*, 21(3): 359-380.
- Ellis, Randall P., Bruno Martins, and Wenjia Zhu. 2016. "Demand Elasticities and Service Selection Incentives among Competing Private Health Plans." Boston University Working Paper.
- Ellis, Randall P. and Wenjia Zhu. 2016. "Health Plan Type Variations in Spells of Health Care Treatment." *American Journal of Health Economics*, 2(4): 399-430.
- Federal Register. 2013. "Medicare and Medicaid Programs: Home Health Prospective Payment System Rate Update for CY 2014, Home Health Quality Reporting Requirements, and Cost Allocation of Home Health Survey Expenses." 78 FR 72255 (2 December 2013): 72255-72320.
- Frank Richard G, and Thomas G. McGuire. 2000. "Mental health." In: Culyer A, Newhouse J, editors. *Handbook of health economics*. Amsterdam: North Holland: 893–954.
- Fronstin, Paul; Martín J. Sepúlveda; M Christopher Roebuck. 2013. "Consumer-Directed Health Plans Reduce the Long-Term Use of Outpatient Physician Visits and Prescription Drugs." *Health Affairs*, 32(6): 1126-1134.
- Haviland, Amelia M., Matthew D. Eisenberg, Ateev Mehrotra, Peter J. Huckfeldt, and Neeraj Sood. 2015. "Do "Consumer-Directed" Health Plans Bend the Cost Curve Over Time?" NBER working paper series 21031, <http://www.nber.org/papers/w21031>.
- Keeler, Emmett B., Willard G. Manning, and Kenneth B. Wells. 1988. "The Demand for Episodes of Mental Health Services." *Journal of Health Economics*, 7(4): 369-392.

- Keeler, Emmett B. and John E. Rolph. 1988. "The Demand for Episodes of Treatment in the Health Insurance Experiment." *Journal of Health Economics*, 7(4): 337-367.
- Kowalski, Amanda. 2016. "Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care." *Journal of Business & Economic Statistics*, 34(1): 107-117.
- Luo, Siyi, and Wenjia Zhu and Randall P. Ellis. 2016. "An Iterative Approach to Estimation with Multiple High-Dimensional Fixed Effects: Controlling Simultaneously for Patients, Providers and Counties." Boston University working paper.
- Manning, Willard G, Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz. 1987. "Health Insurance and the Demand for Medical Care: Evidence from A Randomized Experiment." *American Economic Review*, 77(3): 251-277.
- McGuire, Thomas G. 2016. "Health Insurance for Mental Health Care in Markets with 'Managed Competition' Among Insurers." Working paper.
- Meyerhoefer, Chad D. and Samuel H. Zuvekas. 2010. "New Estimates of the Demand for Physical and Mental Health Treatment." *Health Economics*, 19(3): 297-315.
- Meyerhoefer, Chad D., Samuel H. Zuvekas, and Richard Manski. 2014. "The Demand for Preventive and Restorative Dental Services." *Health Economics*, 23(1): 14-32.
- Newhouse, Joseph P., Charles E. Phelps and M. Susan Marquis. 1979. "On Having Your Cake and Eating It Too: Econometric Problems in Estimating the Demand for Health Services." Santa Monica, CA: RAND Corporation.
- Peters, Rebecca and John Holahan. 2014. "Narrow Networks, Access to Hospitals and Premiums: An Analysis of Marketplace Products in Six Cities." The Urban Institute Research Report, http://www.urban.org/research/publication/narrow-networks-access-hospitals-and-premiums-analysis-marketplace-products-six-cities/view/full_report. (Accessed 7/23/2015)
- Scoggins, John. F. and Daniel A. Weinberg. 2016. "Healthcare Coinsurance Elasticity Coefficient Estimation Using Monthly Cross-sectional, Time-series Claims Data." *Health Economics*, doi: 10.1002/hec.3341.

For Online Publication

Appendix

Table A-1

Plan types market shares

Year	EPO	HMO	POS	PPO	Comprehensive	CDHP	HDHP
2008	2.56	13	14.59	62.81	1.09	5.82	0.13
2009	1.7	19.61	8.95	60.89	2.65	4.93	1.27
2010	0	12.45	9.53	68.78	1.75	6.11	1.38
2011	0.03	10.32	15.54	67.58	1.73	4.52	0.3
2012	0	11.93	18.66	55.97	1.61	11.69	0.15
2013	0	13.3	3.59	75.69	0.37	2.12	4.92
2014	0	6.66	4.28	47.75	2.21	31.47	7.63
All	0.67	12.7	11.54	63.04	1.68	8.62	1.75

Table A-2
OLS estimates

<i>Model: log(Y+1)</i>	Myopic cost share			Forward-looking cost share		
	Coeff.	S.E.	Elasticity	Coeff.	S.E.	Elasticity
All	0.22	(0.006)	0.08	-1.27	(0.02)	-0.41
Outpatient	0.29	(0.009)	0.10	-0.15	(0.01)	-0.05
Inpatient	3.84	(0.096)	0.52	1.58	(0.1)	0.21
Pharmacy	0.31	(0.007)	0.14	-1.75	(0.01)	-0.79
Primary care	0.39	(0.010)	0.12	0.30	(0.01)	0.08
Surgery	1.65	(0.026)	0.52	0.19	(0.03)	0.06
Specialty care	0.16	(0.013)	0.07	0.25	(0.02)	0.11
Lab tests	3.63	(0.171)	0.22	5.16	(0.53)	0.31
Imaging	0.28	(0.010)	0.12	-0.38	(0.02)	-0.15
MH/SA	1.09	(0.029)	0.37	1.01	(0.04)	0.34
Pharmacy	1.85	(0.051)	0.26	1.30	(0.09)	0.18
Room board	0.53	(0.013)	0.16	0.39	(0.02)	0.11
ER	0.34	(0.008)	0.11	0.24	(0.01)	0.07
other	0.30	(0.006)	0.13	-1.71	(0.01)	-0.76

Note: N = 172,122,188 individual months. Each row is a different regression. Each model includes individual*year fixed effects and 84 monthly time dummies. Cluster corrected standard errors use the employer*year*month*family*plan for clusters.

Table A-3
OLS estimates, by detailed Type of Service (TOS)

<i>Model: log(Y+1)</i>	Myopic cost share			Forward-looking cost share		
	Coeff.	S.E.	Elasticity	Coeff.	S.E.	Elasticity
Room and board	0.27	(0.009)	0.09	0.50	(0.02)	0.16
Other inpatient procedures	0.15	(0.034)	0.03	0.51	(0.07)	0.11
Other outpatient procedures	0.62	(0.062)	0.04	0.38	(0.09)	0.02
Mental health	1.65	(0.026)	0.52	0.19	(0.03)	0.06
Substance abuse	0.17	(0.013)	0.07	0.29	(0.02)	0.12
Emergency room	0.37	(0.026)	0.14	-0.13	(0.04)	-0.05
Specialty visits	2.99	(0.149)	0.16	4.98	(0.52)	0.27
Non-specialty visits	2.60	(0.113)	0.17	2.78	(0.41)	0.18
Other	2.04	(0.049)	0.32	1.46	(0.08)	0.23
Other diagnostic services	0.32	(0.011)	0.13	-0.28	(0.02)	-0.11
Dialysis	-1.86	(0.114)	-0.26	-0.88	(0.12)	-0.12
Durable medical equipment	-0.03	(0.011)	-0.01	0.18	(0.03)	0.08
Home visit	0.00	(0.010)	0.00	0.23	(0.03)	0.15
Pharmacy	1.49	(0.905)	0.12	-1.36	(1.03)	-0.11
Specialty drugs	1.09	(0.029)	0.37	1.01	(0.04)	0.34
Transportation	2.20	(0.065)	0.26	2.03	(0.18)	0.24
Laboratory other	1.57	(0.052)	0.24	1.48	(0.12)	0.22
CAT scan	0.96	(0.026)	0.29	1.22	(0.04)	0.37
Mammograms	0.64	(0.033)	0.07	0.49	(0.02)	0.05
MRIS	1.40	(0.039)	0.43	1.69	(0.05)	0.51
PET Scans	2.04	(0.092)	0.31	3.27	(0.2)	0.50
Radiology (x-ray, therapeutic, other)	0.41	(0.013)	0.15	0.45	(0.02)	0.16
Ultrasounds	-0.66	(0.118)	-0.09	0.78	(0.26)	0.11
Surgery	0.71	(0.021)	0.26	0.76	(0.03)	0.27
LTC	0.50	(0.015)	0.14	0.49	(0.02)	0.13
Chiropractic	0.33	(0.009)	0.11	0.29	(0.01)	0.08
Prevention	0.31	(0.007)	0.14	-1.75	(0.01)	-0.79

Note: N = 172,122,188 individual months. Each row is a different regression. Each model includes individual*year fixed effects and 84 monthly time dummies. Cluster corrected standard errors use the employer*year*month*family*plan for clusters.

Table A-4

Detailed estimates in sensitivity analysis

Cross price elasticities using myopic prices				
Model	Outpatient cost share			
	Coeff	SE	Mean CS	Elasticity
Outpatient	-0.37	(0.05)	0.33	-0.12
Inpatient	-0.0004	(0.00)	0.33	-0.0001
Pharmacy	0.03	(0.03)	0.33	0.01
Model	Inpatient cost share			
	Coeff	SE	Mean CS	Elasticity
Outpatient	-24.54	(6.76)	0.14	-3.33
Inpatient	-1.90	(0.33)	0.14	-0.26
Pharmacy	-19.66	(3.68)	0.14	-2.66
Model	Pharmacy cost share			
	Coeff	SE	Mean CS	Elasticity
Outpatient	-0.391	(0.065)	0.44	-0.17
Inpatient	0.0071	(0.003)	0.44	0.003
Pharmacy	-0.627	(0.045)	0.44	-0.28
January and December dropped				
Model	Myopic cost share			
	Coeff	SE	Mean CS	Elasticity
ALL	-0.91	(0.06)	0.35	-0.32
Outpatient	-0.59	(0.05)	0.34	-0.20
Inpatient	-2.45	(0.35)	0.14	-0.33
Pharmacy	-0.64	(0.03)	0.44	-0.28
Alternative IV: total average actual cost share				
Model	Myopic cost share			
	Coeff	SE	Mean CS	Elasticity
Outpatient	-0.78	(0.040)	0.33	-0.26
Inpatient	-1.60	(0.223)	0.14	-0.22
Pharmacy	-0.86	(0.056)	0.44	-0.38

Alternative log transformations

Model	Log(Y+0.1)				Log(Y+10)			
	Coeff	SE	Mean CS	Elasticity	Coeff	SE	Mean CS	Elasticity
ALL	-1.24	(0.066)	0.35	-0.43	-0.64	(0.031)	0.35	-0.22
Outpatient	-0.84	(0.051)	0.33	-0.28	-0.39	(0.022)	0.33	-0.13
Inpatient	-2.16	(0.271)	0.14	-0.29	-1.31	(0.163)	0.14	-0.18
Pharmacy	-0.98	(0.055)	0.44	-0.43	-0.47	(0.022)	0.44	-0.21