The impact of patient cost-sharing on low-income populations: Evidence from Massachusetts[☆]

Amitabh Chandra[a,d], Jonathan Gruber[b,d], Robin McKnight[c,d,*]

[a] Harvard Kennedy School, Harvard University, United States
[b] Department of Economics, MIT, United States
[c] Department of Economics, Wellesley College, United States
[d] NBER, United States

ARTICLE INFO

ABSTRACT

Greater patient cost-sharing could help reduce the fiscal pressures associated with insurance expansion by reducing the scope for moral hazard. But it is possible that low-income recipients are unable to cut back on utilization wisely and that, as a result, higher cost-sharing will lead to worse health and higher downstream costs through increased use of inpatient and outpatient care. We use exogenous variation in the copayments faced by low-income enrollees in the Massachusetts Commonwealth Care program to study these effects. We estimate separate price elasticities of demand by type of service. Overall, we find price elasticities of about −0.16 for this low-income population — similar to elasticities calculated for higher-income populations in other settings. These elasticities are somewhat smaller for the chronically sick, especially for those with asthma, diabetes, and high cholesterol. These lower elasticities are attributable to lower responsiveness to prices across all categories of service, and to some statistically insignificant increases in inpatient care.

© 2013 Elsevier B.V. All rights reserved.

The recently enacted Patient Protection and Affordable Care Act (PPACA) includes the largest expansion of health insurance coverage to low-income populations in our nation's history. The Federal government will spend over $1 trillion over the next decade to subsidize insurance for those below 400% of the Federal Poverty Line (FPL) (Congressional Budget Office, 2013). Roughly, half that total will be through expansions of the Medicaid program, which will provide publicly financed health care for those below 133% of the poverty line at essentially zero patient cost. The other half will be in the form of subsidies to private insurance for those between 133 and 400 percent of the poverty line. These subsidies are of two types: the first type is premium subsidies, which offset the premium cost of insurance by limiting the percentage of income that low-income individuals must pay. The second type is cost-sharing subsidies, which offset to some extent the copayments, coinsurance and deductibles that these low-income populations face.

The motivation for the subsidies is twofold: to make the transfers in PPACA more progressive, and to protect low-income populations from sacrificing necessary medical care because of cost. The optimal level of such subsidies, therefore, depends critically on the way in which the medical care utilization of low-income groups responds to cost sharing, and how any change in utilization impacts their health. On one hand, greater patient cost-sharing could help reduce the fiscal pressures associated with insurance expansion by reducing the scope for moral hazard. But on the other hand, there has been speculation that low-income patients may be more price sensitive than other patients or that low-income patients may be more likely to experience adverse health consequences as a result of cost-sharing (Baicker and Goldman, 2011).

Differential effects on low income patients could arise for a number of reasons. First, low-income patients may simply be more responsive because they face a tighter budget constraint; in this case, we would expect low-income patients to cut back on care with the lowest marginal benefit. Second, it is possible that lower income individuals are less able to evaluate the marginal benefit of their care than higher income individuals and, as a result, may have a higher propensity to cut back on high marginal benefit care. In their study of drug copayments in Medicaid, for example, Reeder and Nelson (1985) argued that, because education is positively correlated with income, low-income individuals may be less able to communicate with their physicians and, consequently, make less

well-informed decisions. Goldman and Smith (2002) provide evidence that patients of lower socioeconomic status are less likely to adhere to treatment regimens for chronic illnesses and, as a result, experience worse health outcomes. Third, higher rates of chronic illness among low-income populations could imply differential effects of cost-sharing in low-income populations. In prior studies that have identified adverse consequences, these consequences have generally been concentrated among the chronically ill (Goldman et al., 2007).

These concerns highlight the importance of considering the possibility of cross-price effects in response to copayment increases: cost-sharing for one service could reduce utilization of complementary services or it could increase utilization of substitute services. The positive cross-price effects on substitute services are known in the empirical literature as "offset effects," because decreases in one service are offset by increases in substitute services (McGuire, 2012). In particular, prior work has suggested that reduced primary care due to copayment increases can lead to higher downstream costs in inpatient or outpatient settings. Commentators such as Evans et al. (1993) have argued that user fees are a "policy zombie", while Donaldson (2008) argued that it is "wrong, unfair, and ineffective to try to limit consumer and patient access through user fees, and also to dress up this process as actually enhancing access."

In this paper, we use exogenous variation in the copayments faced by low-income enrollees in Massachusetts' Commonwealth Care program. This state program was the model for PPACA, providing highly subsidized insurance for families below 300% of the FPL. Importantly, there was a substantial increase in the copayments paid by enrollees in this program in July 2008 that we exploit for identification. Using unique claims data on the Commonwealth Care population provided to us by the state, we are able to estimate the effect of greater cost-sharing on overall utilization. Because copayments increased nearly proportionately across most services, we are unable to separately estimate own- and cross-price elasticities for most services. However, we are able to estimate the overall price elasticity associated with across-the-board copayment increases. In addition, because copayments for inpatient admissions did not increase for most of our sample, we are able to provide evidence on hospitalization offsets.

In previous work (Chandra et al., 2010b), we had provided preliminary evidence on the price elasticity of demand in this population, using only the first six months of post-policy change data. This paper extends that analysis by incorporating additional post-policy change data, by using a superior estimation framework, and most importantly by considering a variety of questions not addressed in that earlier paper, such as the issue of population heterogeneity in demand elasticities and in offset effects. We are also able to perform a number of specification checks that provide assurance that our results are not be driven by stockpiling (where patients avail of medical care immediately before the policy change and do not use care in the immediate months after), or other confounders that may be timed with the policy change. As a result of a variety of improvements to our data and estimation, our current estimate for the overall price elasticity of demand in this population is on the low end of the range of elasticities that we reported in our earlier work. Overall, we find price elasticities of about −0.16 for this low-income population, which is similar to, but somewhat lower than, elasticities calculated for higher-income populations in other settings.

We also find lower price elasticities among individuals with chronic illness and with higher levels of prior spending, suggesting that copayments are less important in these subsamples. In addition, we find no evidence of offsetting increases in hospitalizations in response to the higher copayments, although there are some statistically insignificant impacts among the chronically ill population.

Our paper proceeds as follows. Section 1 reviews the prior literature on this topic. Section 2 provides background on the institutional setting and data. Section 3 presents our estimation strategy. Section 4 shows our results, while Section 5 presents a set of robustness checks of those findings. Section 6 concludes with a discussion of the implications of our findings.

1. Prior literature

Most analyses of price sensitivity of medical care still rely on the evidence from the RAND Health Insurance Experiment (HIE) of the mid-1970s (Newhouse, 1993). This seminal study found that, for the population as a whole, there was a significant but modest response of medical care utilization to the point-of-service cost of medical care. Notably, the reductions in care appeared to come across the board, both in categories of "effective" and "ineffective" medical care. There were, however, no "offset" effects in terms of reduced primary care leading to a demand for more hospital care; indeed, reduced primary care appeared to lower spending on hospital care. Most importantly, there was no evidence of a detrimental effect on the typical person in terms of worsened health. That is, for the average person in the HIE, there did not appear to be productive returns to marginal health care utilization in terms of improving health status.

Although the sample size was more limited, the HIE did consider heterogeneity by income and health. For the subset of low-income population, the findings paralleled those for the larger population: a modest impact on health care spending, with no offset effects and no impact on health status. However, there was an important potential exception to this finding: for the chronically ill low-income population, there was a suggestion of a sizeable rise in blood pressure for those in the higher cost-sharing insurance plan (although the overall health results for this population were mixed, and the blood pressure result itself was not significant). But the HIE evidence is now over thirty years old, and changes in the practice of medicine—including greater reliance on managed care contracts, the availability of many new pharmaceutical and surgical interventions, the growth of imaging and diagnostic technology, and the development of the medical device industry—may imply a structural change in the elasticity of medical demand and the health impacts of any utilization reductions.

Besides the HIE, there has been relatively little research on the impacts of cost-sharing on low-income populations. Indeed, Baicker and Goldman (2011) write that "while there is a lot of speculation that the poor have more elastic demand, there is little evidence." Much of the existing work focuses on the introduction of copayments for prescription drugs in Medicaid programs. Nelson et al. (1984) examined the introduction of a drug copayment in South Carolina's Medicaid program and found that it was associated with a statistically significant decline in drug purchases, in a pre-post framework and in comparison with Tennessee's Medicaid program. Follow-up work in Reeder and Nelson (1985) examines the change in drug utilization within South Carolina and notes a reduction in classes of drugs which, if not taken, "could result in a deterioration of health that could ultimately lead to the use of more expensive medical services." Stuart and Zacker (1999) use cross-state variation in the use of drug copayments in Medicaid to examine the impact of copayments on Medicare dual eligibles. They find that individuals in states with copayments, ranging from $0.50 to $3.00, use 15 percent fewer prescriptions than individuals in states without copayments, with the impact resulting primarily from a reduction in drug use on the extensive margin.

Cunningham (2002) analyzes cross-state variation in Medicaid's use of prescription drug cost-control methods, including copayments, and finds that cost-control efforts are associated with reductions in reported access to prescription drugs. Our research builds on this existing evidence, by exploiting plausibly exogenous variation affecting the low-income population that allows us to calculate a price elasticity of demand for this population.

The theoretical impact of cost-sharing on overall spending is complicated by cross-price effects (Baicker and Goldman, 2011). On one hand, cross-price effects could imply that a copayment for one service would reduce utilization of complementary services; an office visit copayment, for example, might reduce spending on labs and diagnostic testing. On the other hand, cross-price effects could imply that a copayment for one service would increase utilization of substitute services; reduced use of primary care can lead to worsened health, which might indirectly lead to increased utilization of emergency or inpatient services. Offset effects are a particular concern if psychological biases cause patients to overweight the immediate cost of a copayment relative to expected future health benefits (Baicker et al., 2012). In the potential presence of cross-price effects, it is important to evaluate empirically the impact of cost-sharing on overall spending and on indicators of worsening health, such as hospitalizations.

The possibility of offset effects is real: while the RAND Health Insurance Experiment did not find any evidence for offsets, more recent research has suggested that increased cost-sharing may be associated with offsetting increases in other forms of care. For example, Gaynor et al. (2007) find that increased spending on outpatient care offset 35% of the savings from reduced prescription drug use after drug copayment increases in employer-provided health insurance. Such offsets are often found to be concentrated among chronically ill populations. For example, in our own earlier work (Chandra et al., 2010a), we found evidence that increased prescription drug and office visit copayments caused offsetting increases in hospitalizations in elderly who had been diagnosed with chronic diseases such as diabetes, hypertension, and hyperlipidemia. Similarly, Trivedi et al. (2010) find that increased copayments for ambulatory care for elderly patients were associated with increased rates of hospitalizations; these effects were concentrated among enrollees of lower socioeconomic status and enrollees who had chronic illness or a history of myocardial infarction. In a paper that is especially relevant to our setting, Tamblyn et al. (2001) examined the effect of cost-sharing on prescription drugs in the poor and elderly and found an increase in the probability of hospitalization, nursing home admission, or mortality. The key role of cost-sharing in the PPACA suggests the importance of revisiting this topic in a low-income population.

2. Institutional setting and data

Our study examines the impact of increased patient copayments that were imposed on low-income adults in the Commonwealth of Massachusetts. As part of the insurance expansion that was signed into law on April 12, 2006, the state of Massachusetts provides completely subsidized coverage to legal residents earning up to 150% of the FPL and highly subsidized coverage to those earning between 150% and 300% of the FPL.[1] These subsidies are provided through a program known as Commonwealth Care (henceforth, CommCare) to persons between the ages of 19 and 64 who are not offered employer-provided health insurance, cannot be dependents on another person's family plan, and do not qualify for other public health insurance programs such as Medicaid.

Members are eligible for different plan types based on where their income is relative to the FPL, and each plan type offers a different level of patient cost-sharing. At the beginning of our sample period in July 2007, there were four different plans— each with the same benefits but with different levels of patient cost-sharing. None of the plans had deductibles, so all the cost-sharing took the form of copayments. Plan 1 was for all members with incomes below 100 percent of the FPL and it had the lowest level of cost-sharing ($0 for all services, except emergency room visits and prescription drugs). Plan 2 was for all members with incomes between 100 percent and 200 percent of the FPL and it had higher copayments for all services than did Plan 1. This plan structure generated a discontinuity in cost-sharing at 100 percent of the FPL. Members with incomes between 200 and 300 percent of the FPL had a choice of two plans, with different levels of cost-sharing and premiums. If they chose Plan 3, they had higher cost-sharing than Plan 2; if they chose Plan 4, they had the same cost-sharing as member of Plan 2, but they had to pay a higher premium than for Plan 3. Thus, there was a cost-sharing discontinuity at 200 percent of the FPL, but only for members who opted for Plan 3.

In July 2008, copayments were increased for Plans 2 and 3, generating changes in the discontinuities at 100 percent of the FPL and 200 percent of the FPL. In addition, Plan 4 was discontinued in July 2008 and its members were placed into Plan 3.[2] Consequently, Plan 4 enrollees experienced a substantial increase in patient cost-sharing as a result of the elimination of Plan 4. These changes in the discontinuities in copayments at 100 and 200 percent of the FPL that occurred in July 2008 are the key source of variation that is exploited in our empirical work below.

CommCare charges no monthly premium for members with incomes below 150% of the FPL. For those with incomes above 150% of the FPL, there is a monthly premium, which increases with income. Currently, the monthly premium for the lowest cost plan is $40 for enrollees with income between 150 and 200 percent of the FPL, $78 for those between 200 and 250 percent of the FPL, and $118 for those between 250 and 300 percent of the FPL.

2.1. Data

For the purposes of this analysis, Massachusetts' Commonwealth Health Insurance Connector Authority provided us with the universe of (de-identified) enrollment and claims data from July 2007 through June 2009. This sample period covers a full year before and a full year after the copayment change that we are studying. From the enrollment file, we can observe the CommCare plan in which each individual was enrolled in each month of the sample period. We also observe very basic demographic information, such as age and gender. For most individuals, we observe household income as a share of the FPL, although this information is missing for some individuals. As a result of this missing information, we must exclude 25% of member-months from our sample. Because income information is disproportionately missing for individuals who were enrolled for shorter periods of time, this restriction affected only 16% of the member-months in the continuously enrolled subsample that is the focus of our analysis. We also exclude 10,073 member-months that report income levels above 300% of FPL and 19,290 member-months that report ages below 19 or above 65, since these individuals should not

---

Table 1
Member characteristics.

	All members	Continuously enrolled (no plan switching)
Age	40.3	40.9
Percent Male	47	47
Percent enrolled in Plan 1	52	52
Percent enrolled in Plan 2	36	38
Percent enrolled in Plan 3	10	9
Percent enrolled in Plan 4	1	1
Percent with Charlson score of 1 or more	23	24
Percent with Chronic disease	29	31
Average monthly expenditure		
Total	$384.6	$358.8
Hospitalizations	$86.2	$77.7
ER	$43.2	$39.7
Drugs	$44.5	$43.9
Office visits	$46.8	$44.5
Outpatient	$79.7	$75.4
Lab	$58.2	$54.5
Average FPL	95.4%	96.1%
Rate of avoidable hospitalizations	0.06%	0.05%
Number of members	247,565	122,456
Number of member-months	2,842,493	1,865,777

be eligible for CommCare. After making these adjustments, our full sample includes 2,842,493 member-month observations and our continuously enrolled subsample includes 1,865,777 member-month observations.

Our claims data include the universe of medical claims for CommCare enrollees during our sample period. For each member in each month of our data, we use the claims data to calculate the sum of all spending for each type of medical care, including zeroes for members who have no claims in a particular month.[3] For our regression analysis, we then collapse these member-month observations to averages for each percent of income as a share of the FPL in each of the 24 months of data. We allow two observations per month for income cells between 200 and 300 percent of FPL: one for individuals who were ever enrolled in Plan 4 and one for individuals who were never enrolled in Plan 4. As a result, the sample size for our main regressions is 9,644.

2.2. Descriptive statistics

In Table 1, we describe the characteristics of our study population and, in Table 2, we provide details of the copayment change. As Table 1 illustrates, the sample is poor; on average, members were at 95% of the FPL. They are also quite sick, with about 30% of the sample having some form of chronic disease (as measured by a diagnosis of diabetes, hypertension, hyperlipidemia, asthma, arthritis, affective orders, or gastritis— conditions that are derived from Goldman et al. (2004) and also used in Chandra et al. (2010a)). Annual spending is approximately $4,600, much of which comes from hospitalizations and outpatient care.

In our data, we identify outpatient care as care that indicates "outpatient hospital" as the place of service on the claim; it includes both facility and professional claims for the outpatient hospital visit. Such care accounts for about 20% of total costs. Labs include any claims that indicate "Independent Laboratory" as the place of

service, or that indicates that the service provided was "Laboratory" or "Radiology." Labs account for approximately 15% of total costs. Office visits include professional claims that do not indicate that the place of service is outpatient hospital, independent laboratory, or emergency room; that do not occur on the same day as any inpatient hospital claims; and that do not have a service code of "Laboratory" or "Radiology." They are, therefore, largely ambulatory visits for checkups and consultations and they account for roughly 15% of total costs.

In Table 1, we report means for two samples: the full sample with 247,565 members and a sample of 122,456 members who stayed continuously enrolled at the time of the copayment increase. Individuals in this second "continuously enrolled" sample were not necessarily enrolled for the full 24-month sample period, but did not change their enrollment status at the time of the copayment increase. Individuals who initially enrolled in Plan 4 but switched to Plan 3 in July 2008 are not excluded from the "continuously enrolled" sample, since it was not possible for them to remain in Plan 4 at that time.

We created this "continuously enrolled" sample because examining utilization changes in the full sample is subject to considerable bias if member income or enrollment is a function of health status. This is a real concern in the full sample: in our work on the individual mandate in Massachusetts, we demonstrated that the insurance market for Commonwealth Care was laced with adverse selection prior to the individual mandate binding on January 1, 2008 (Chandra et al., 2011). Consistent with adverse selection, we found that those who enrolled first in CommCare were sicker, with relatively healthy individuals joining the plan closer to the date of the mandate. This time-varying change in the composition of enrollees will affect full-sample estimates of the price elasticity of demand because utilization and health-status are correlated in a time varying manner. Moreover, there was a change in the premiums charged to enrollees in CommCare at the same time that copayments changed, which could impact the composition of enrollees. To mitigate these concerns, we examine the continuously enrolled sample, who were continuously enrolled just before and after the policy change, to ensure that our findings are not the result of differential entry and exit rates for different types of individuals around the time of the copayment change. Comparing the means in Table 1 reassures us—the two samples are similar in age, gender, and income; furthermore, utilization levels

---

[3] In our earlier work, we computed average monthly utilization for each individual over the months during which they were enrolled and then computed an average across all individual-years. In this version, we simply compute the average of monthly utilization across all member-month observations. In effect, we previously placed disproportionate weight on individuals who were enrolled for shorter periods of time; now the weight that we place on individuals is proportional to the amount of time that they are enrolled.

Table 2
Copayment changes by percent of federal poverty line of member.

	Plan 1: 0–100 percent of FPL	Plan 2: 100–200 percent of FPL	Plan 3: 200–300 percent of FPL	Plan 4: 200–300 percent of FPL
Copayments for hospital inpatient admission				
October 2006–June 2008	$0	$50	$250	$50
July 2008–present	$0	$50	$250	$250
Copayments for emergency room visits				
October 2006–June 2008	$3	$50	$75	$50
July 2008–present	$3	$50	$100	$100
Copayments for outpatient surgery				
October 2006–June 2008	$0	$50	$100	$50
July 2008–present	$0	$50	$125	$125
Copayments for office visits (primary care/specialist)				
October 2006–June 2008	$0/$0	$5/$10	$10/$20	$5/$10
July 2008–present	$0/$0	$10/$18	$15/$22	$15/$22
Copayment for prescription drugs (generics/formulary/non-formulary retail)				
October 2006–June 2008	$1/$3	$5/$10/$30	$10/$20/$40	$5/$10/$30
July 2008–present	$1/$3	$10/$20/$40	$12.50/$25/$50	$12.50/$25/$50
Average copayment change (across all services)				
October 2006–June 2008	$1.11	$ 8.39	$16.83	$ 8.39
July 2008–present	$1.11	$13.99	$21.30	$21.30
Increase	0 percent	51 percent	24 percent	93 percent

Note: Prior to July 2008, there were two plan types for members whose income was between 200 and 300 percent of the FPL (Plans 3 and 4). Plan 4 was discontinued in July 2008 and its members were enrolled in Plan 3. Throughout this period, the copayment for emergency room use is waived if the patient is admitted to the hospital. To calculate average copayments, we weighted copayments for each service (hospitalizations, ER use, prescription drugs) by its pre-reform utilization (pre July 2008 utilization). Source: The Massachusetts Health Insurance Connector Authority (2008).

are similar as are enrollments across the four plans that we study.[4] Nonetheless, we acknowledge the possibility that our findings might not generalize to the broader, low-income population.

In Table 2 we summarize the variation that we use to understand the effect of patient cost-sharing on patient demand: we exploit a series of changes in the discontinuities in the cost-sharing provisions of CommCare that occurred in July 2008. These changes, which varied by type of plan, are reported in Table 2. Of these changes, the most salient are the increases in copayments for prescription drugs and office visits for members whose incomes exceeded 100% of the FPL. Copayments for ER visits increased only for those with incomes over 200% of FPL. Copayments for hospital admissions were unchanged, except for the Plan 4 members who had to switch to Plan 3. Note that there were no copayment changes for members whose income was between 0 and 100 percent of the FPL.

In order to estimate the price elasticity of demand, we devise a simple measure of the overall change in copayments. Specifically, we calculate the weighted average of the copayments for generic drugs, formulary drugs, non-formulary drugs, office visits, ER visits, and hospitalizations. The weights are average monthly utilization of each type of care during the pre-period (July 2007 to June 2008) for the entire sample. At the bottom of Table 2, we summarize the (unadjusted) change in copayments that we will exploit: that is, there is no change for Plan 1, a 51% increase for Plan 2, a 24% increase for Plan 3, and a 93% increase for Plan 4. This weighted average measure produces meaningful results in two situations: One such situation is when copayment changes are proportional across all services. Another is if individuals respond to copayment increases across services in a manner that is proportional to the overall
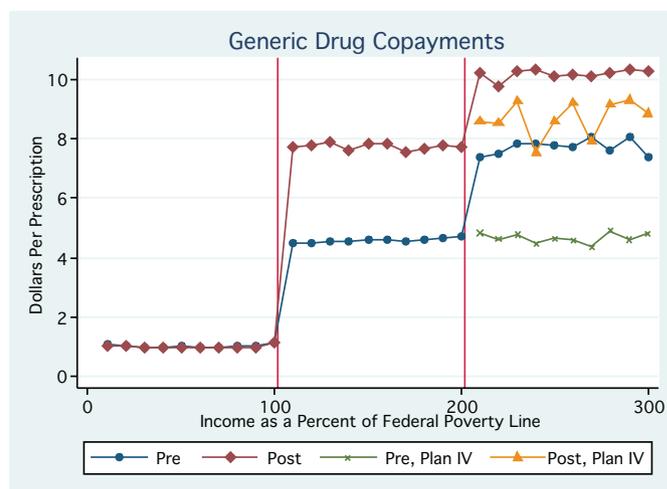

Fig. 1. Copayments for generic drugs.

frequency that those services are used i.e., more weight is assigned to office visit and generic prescription drug copayments, because those are the most frequently used in the population overall.[5]

In Fig. 1, we plot the average copayment for one category of utilization—generic drugs—by income category and by period ('pre' denotes the period before July 2008 and 'post' denotes the period after July 2008). This figure graphically illustrates the basis for our formal regression analysis below.

---

[4] We investigated the possibility that a disproportionate number of people switched plans because of the copayment increase. The percent of members who were new members of a plan in May, June, July and August of 2008 was 5.2, 5.2, 4.7 and 4.4 percent respectively (the policy change occurred in July 2008). Similarly, the percent of members who left the plan in May, June, July and August of 2008 was 5.2, 4.8, 4.7 and 5.7 percent respectively. In other words, we do not see excess churning at the time that the policy change occurred.

[5] It is also possible that individuals respond more to the copayments that, proportionately, have affected them more in the past. To test sensitivity to this issue, we calculated an alternative copayment measure, with weights that are specific to an individual's spending tercile in the base year. When we calculate the weighted average copayment in this way (and control for tercile in our regressions), our overall estimates were very similar to those reported in the paper, suggesting that this possibility, while plausible, was not enough to justify the alternative approach.

Table 3
GLM estimates of the effect of copayment increases on healthcare spending (all patients).

	(1) Total spending	(2) Hospital spending	(3) ER spending	(4) Outpatient spending	(5) Office visit spending	(6) Rx spending	(7) Lab spending
Panel A: all patients							
ln(Copay)	−0.158 (0.064)	−0.115 (0.247)	−0.207 (0.152)	−0.203 (0.122)	−0.147 (0.051)	−0.131 (0.074)	−0.282 (0.089)
N	9644	9644	9644	9644	9644	9644	9644

Note: Each number in the table is an elasticity that comes from a separate GLM regression. Dependent variable is Total Monthly Healthcare Expenditure per member per month. Standard errors are clustered on member, and explanatory variables include percent of FPL fixed-effects, and year*month indicators.

Because a number of cost-sharing increases occurred at the same time, it is important not to attribute changes in a specific category of spending, such as the utilization of prescription drugs, solely to changes in the copayments for these drugs. Declines in the use of prescription drugs may, for example, partially reflect a reduction in the use of office visits. For this reason, we are most interested in the reduction in total expenses across all categories of care.

3. Estimation

Because the policy variation occurs at the level of plans, and plans are defined according to the percent of FPL of the member's income, we estimate our regression models at the level of FPL $\times$ month, and cluster our standard errors at the level of FPL. The simplest framework to think about estimating the price elasticity of demand is with the equation below:

$$\ln Spending_{jt} = \beta_0 + \beta_1 \ln Copayment_{jt} + FPL_j + Time_t + \varepsilon_{jt} \quad (1)$$

Here, $Spending_{jt}$ is the average monthly spending for income cell j in time t, and $Copayment_{jt}$ is the weighted average copayment for this cell at a point in time. Income is defined as percent of the FPL so all our specifications can (and do) include fixed effects for each percent of FPL. Note that plan fixed effects are perfectly correlated with the percent of FPL fixed effects, and are therefore not included. We also include year $\times$ month fixed effects to remove any effects of time trends or seasonality.

Eq. (1) identifies the price (copayment) elasticity of healthcare as β_1. In Eq. (1) variation in copayments is coming only from within-plan increases. Our approach is identical to identifying the effect of copayment changes using a regression-discontinuity design, where we identify off changes in copayments at 100 and 200 percent of the poverty line (see Fig. 1); later in this paper (Table 7), we restrict the analysis to utilization in the period immediately around the policy change and demonstrate that the results are unchanged.

We could estimate Eq. (1) with OLS, but this would make sense only in a world where all cells had non-zero spending. It is tempting to add a small amount of (arbitrary) spending, such as $1, to overcome the problem that ln(0) is not defined; indeed, this is the approach that we took in Chandra et al. (2010b). But this can be a problematic strategy, because ln(Spending) is quite sensitive to spending at the bottom of the spending distribution; converting ln($1) to ln($2) has the same effect as converting ln($1000) to ln($2000). To circumvent this problem, we follow the literature and estimate a general linear model with a log-link function (Buntin and Zaslavsky, 2004). In GLM models with a log link-function, the conditional mean is modeled as:

$$\ln E(Spending|X) = \beta_0 + \beta_1 \ln Copayment_{jt} + FPL_j + Time_t \quad (2)$$

This setup models both the mean and variance on the original scale of the dependent variable (here, dollars). We allow for a natural form of heteroscedasticity, where $var(Spending|X)$ is allowed to depend on the mean level of conditional spending, $E(Spending|X)$, which is a function of the covariates. The use of the GLM model

avoids the problem with two-part models (which introduce selection on the dependent variable) and yields estimates that are easily interpretable as elasticities.[6] In principal, Eq. (2) could also be estimated at the level of individuals and we could include individual fixed effects. Because the policy change occurs at the plan level, and individuals are assigned to plans based on their income, there is no additional benefit from this approach.

4. Results

Table 3 reports our key results. In column 1, we report the overall elasticity, which is −0.158; a 10 percent increase in prices faced by patients would reduce utilization by 1 to 2 percent. This is comparable to the much older estimates of RAND HIE, albeit somewhat lower.

With our data, we are able to estimate separate effects by type of service: hospital spending, ER spending, outpatient spending, office visit spending, prescription drug spending, and spending on laboratory services. We are also able to examine whether there were 'offset' effects; that is, did inpatient hospital visits increase, potentially because of excessive reductions in primary care utilization?

The results in Table 3 suggest that elasticities for all services fall in a relatively narrow range, between −0.1 and −0.3. Perhaps most interestingly, we find no evidence of offset effects. Hospital utilization falls, rather than rising, as copayments are increased (but this effect is statistically insignificant). For Plans 1, 2 and 3, hospital copayments did not increase, so that any impact on hospital utilization represents a pure offset effect. But for patients in Plan 4, this could reflect the demand side impact of copayments rising from $50 to $250 for hospital care. Our results are similar, however, if we exclude Plan 4 (which is unsurprising since Plan 4 is 1 percent of the total sample), clarifying that we are estimating the offset effect only.

We next examine whether the overall elasticity varies by type of patient. We first examine whether price sensitivity varies with health status, since prior research has suggested that any offset effects are likely to be concentrated among those in poor health (Manning et al., 1987; Chandra et al., 2010a). We use two measures of chronic illness. One measure is an indicator for the presence of a diagnosis of hypertension, high cholesterol, diabetes, asthma, arthritis, affective disorders or gastritis, following definitions in Goldman et al. (2004). Thirty-one percent of our continuously enrolled sample meets this definition of chronic illness. The second measure is the Charlson Comorbidity Index, a weighted sum of indicators for certain health conditions, where the weights indicate the relative increase in one-year mortality risk associated with the condition, as described in Charlson et al. (1987). We report separate price elasticities for individuals with a Charlson index value

---

[6] In contrast to this model, in OLS one is assuming that $E(\ln Spending|X) = \beta_0 + \beta_1 \ln Copayment_{jt} + FPL_j + Time_t$, which is difficult to translate into easily interpretable statements about $E(Spending)$. The solution is to use a 'smearing' estimator, where exponentiated predictions are multiplied by a smearing factor to calculate expected values on the raw (unlogged) scale (Manning et al., 1987).

of 0 (i.e., those who have none of the conditions) and for those with Charlson index values of 1 or greater. Twenty-four percent of our continuously enrolled sample meets this definition of chronic illness. These chronically ill individuals have much higher spending levels than the full sample of continuously enrolled individuals; those with a chronic illness spend, on average, $563 per month and those with a Charlson score of 1 or greater spend, on average, $787 per month, as compared to $359 per month for the entire sample of continuously enrolled individuals.

We also examine whether price sensitivity is different in the "tail" of the spending distribution. To do this, we divide our sample into individuals with high and low levels of spending during a baseline period, the 2007–08 plan year. The "high spenders," those who were in the top quartile of utilization during the base period, spend, on average, $968 per month, whereas those who were in the bottom three quartiles spend, on average, $158 per month.

The results in Table 4 show total spending elasticities for different subsamples of our data. There is a somewhat higher elasticity for those with a Charlson index score of 1 or more, as seen in column 3, although it is not statistically significantly different from the elasticity for those with a Charlson index score of 0. There is a substantially smaller elasticity for those with a chronic illness diagnosis, seen in column 5; this smaller elasticity appears to be driven by positive, but insignificant, relationships between copayments and overall spending among those with diabetes and high cholesterol, and a positive relationship that is statistically significant at the 10% level for asthma patients.[7] The positive, statistically insignificant result for asthma patients is shown in column 6. Similarly, the price elasticity is positive and insignificant for those who were in the top quartile of the spending distribution, as seen in column 7. Taken together, these results suggest that the impact of copayments is not uniform; there appears to be less price sensitivity among less healthy enrollees.

There are two classes of explanations for this set of findings. One is that less healthy patients are less responsive to higher copayments, because the marginal benefit of care to them exceeds the higher out-of-pocket cost of the copayment. The other is the presence of offset effects amongst those who are chronically ill. Indeed, this is the group for whom offsets have most frequently been found in prior work (Chandra et al., 2010a; Goldman et al., 2007; Trivedi et al., 2010). Of course, both explanations could apply.

In Table 5, we explore these explanations by restricting the analysis to members with any chronic illness. We observe smaller elasticities for presciption drug and office visits spending among the chronically ill, consistent with the first explanation. In this sample, we do not see statistically significant evidence of hospital spending increases as a result of copayments increasing, although the point estimate is positive. Findings (not reported) are similar for the members in the top quartile of the baseline spending distribution, who have a statistically insignficant increase in hospital and outpatient spending. These findings are suggestive of offset effects but, given the imprecision of the estimates, are not conclusive. Regardless of the mechanism, the lower price elasticities among these subsamples suggest that, while copayments are effective on average at reducing total spending among the low-income population, there are chronically ill subsamples for whom copayments are less important.

The overall reduction in utilization reflects the fact that some members cut back on services completely (extensive margin), while others reduce their utilization while continuing to use some

[7] High cholesterol and asthma diagnoses receive a score of 0 in the Charlson index, so the impacts of cost-sharing for these diagnoses explains the difference in our findings when we use the chronic disease indicator instead of the Charlson index.

Table 4
GLM estimates of the effect of copayment increases on medical spending (by type of patient).

	(1) All patients	(2) Charlson index = 0	(3) Charlson index = 1 +	(4) No chronic disease	(5) Any chronic disease	(6) Asthma	(7) Top spending quartile	(8) Excluding top quartile
ln(Copay)	−0.158 (0.064)	−0.184 (0.068)	−0.216 (0.100)	−0.274 (0.086)	−0.057 (0.094)	0.366 (0.215)	0.131 (0.083)	−0.161 (0.076)
N	9644	9619	9590	9606	9608	7893	9590	9627

Note: Each number in the table is an elasticity that comes from a separate GLM regression. Dependent variable is Total Monthly Healthcare Expenditure per member per month. Standard errors are clustered on member, and explanatory variables include percent of FPL fixed-effects, year*month indicators.

Table 5
GLM estimates of the effect of copayment increases on spending (chronically ill members).

	(1) Total spending	(2) Hospital spending	(3) ER spending	(4) Outpatient spending	(5) Office visit spending	(6) Rx spending	(7) Lab spending
ln(Copay)	−0.057 (0.094)	0.356 (0.407)	Insufficient data	−0.194 (0.174)	−0.100 (0.060)	−0.098 (0.073)	−0.318 (0.101)
N	9608	9608	9608	9608	9608	9608	9608

Note: Each number in the table is an elasticity that comes from a separate GLM regression. Dependent variable is Total Monthly Healthcare Expenditure per member per month. Standard errors are clustered on member, and explanatory variables include percent of FPL fixed-effects, and year*month indicators.

Table 6
OLS estimates of the effect of copayment increases on any spending (extensive margin results), (all patients).

	(1) Any spending	(2) Any hospital spending	(3) Any ER spending	(4) Any outpatient spending	(5) Any office visit spending	(6) Any Rx spending	(7) Any lab spending
ln(Copay)	−0.052 (0.009)	−0.0003 (0.002)	−0.001 (0.002)	−0.001 (0.005)	−0.019 (0.007)	−0.069 (0.010)	−0.024 (0.005)
N	9644	9644	9644	9644	9644	9644	9944
Average spending|spending > 0	$746	$5605	$1055	$697	$206	$147	$362
Implied elasticity from extentive margin	−0.11	–	–	–	−0.088	−0.231	−0.159

Note: Each number in the table is the percentage point change in the probability of having any medical spending. Each estimate comes from a separate OLS regression that is estimated at the percent of FPL × month level. Dependent variable is the fraction of the cell that had any utilization. Standard errors are clustered on member, and explanatory variables include percent of FPL fixed-effects, and year*month indicators. The implied elasticity from the extensive margin is the price elasticity of demand that would arise from behavior at the extensive margin. Empty cells are those where the effect at the extensive margin is not statistically different than zero.

care (intensive margin). These two effects are combined in the tables that we have reported so far. In Table 6, we report the importance of the extensive margin for all patients. Here, the dependent variable is the fraction of members who had any utilization of care, by type of care; this can modeled by OLS. We see statistically insignificant, negative effects on the extensive margin of hospital care, emergency room care, and outpatient care. But a 10% increase in copayments results in a statistically significant 0.19 percentage point decline in the probability of making any physician office visits, a 0.69 percentage point decline in the use of any prescription drugs, and a 0.24 percentage point decline in the use of any laboratory services. These results are very consistent across samples.

We can ask how much of the overall decline in utilization is plausibly the consequence of patients reducing utilization at the extensive margin. To perform this calculation, we multiply the change in the margin of participation at the extensive margin with the spending among those who adjusted on the extensive margin, and then divide by the total spending. We proxy utilization for those who adjusted on the extensive margin by using the average utilization of patients who had positive utilization.[8] For example, for total spending, we calculate the implied elasticity from behavior at the extensive margin as −0.052 × $746/$358.8 = −0.11 (here, $746 is the average spending of enrollees with non zero spending, and $358.8 is average spending from Table 1). This is two-thirds of the total elasticity reported in Table 3, and suggests that about 70 percent of the decline in spending might be attributed to members dropping to zero utilization; the rest of the decline is a consequence of members using fewer services after copayments increase than before. Table 6 performs the same calculation for different categories of care (we ignore categories where behavior at the extensive margin was not statistically different than zero): in general,

behavior at the extensive margin explains a substantial portion of the overall price elasticity. This finding is consistent with Stuart and Zacker (1999), who concluded, in their cross-sectional analysis of drug copayments in Medicaid, that the primary impact of drug copayments in Medicaid was a reduction in Medicaid recipients filling any prescriptions. It is a significant finding for it suggests that rather than all patients cutting back on some care, some patients cut back completely.

5. Robustness checks

In Table 7 we examine the robustness of our results to three alternative specifications. In the first specification, we ignore information from the window just before and just after the policy change (that is, we exclude June and July, 2008). We did this to examine the concern that patients may be stockpiling drugs, or seeking care, in anticipation of the policy change. Such behavior will cause us to overstate the price elasticity of demand for it accentuates the reduction in utilization after the copayment increase (the increase in pre-policy use is accompanied by a reduction in post-policy use). The first panel of Table 7 shows that this fear is unfounded: the elasticities reported here are quite similar to those reported in Table 3, where we use all the data available to us. The use of prescription drugs is particularly susceptible to the concern about stockpiling, but we do not find any support for this actually happening.

Next, we restrict the analysis to the period just before and just after the policy change (that is, we keep data from April through September of 2008). The motivation for this specification is that the US financial markets collapsed in early October of 2008 and there was a rapid rise in recession-induced unemployment in late 2008 and 2009. We wanted to be sure that the price elasticity of demand that we are estimating was not partially confounded by these trends in the business cycle. While it is true that we have included income fixed-effects and very flexible year × month controls, there may still be income-specific trends in the data. To allay this concern we constructed Panel B of Table 7. The results are similar to those in Table 3, so we are reassured that the business-cycle is not an alternative explanation for our results. One downside of

---

[8] Note that if all members had positive spending pre-reform, and if (pre-reform) spending for the average member were the same as spending for a person who dropped to zero utilization, then this exercise would generate the same elasticity as what we have estimated above. Of course, we have no way of validating this assumption, so this calculation is purely illustrative of the potential importance of these extensive margin responses.

Table 7
Specification and falsification checks: GLM estimates of the effect of copayment increases on spending (all patients).

	(1) Total spending	(2) Hospital spending	(3) ER spending	(4) Outpatient spending	(5) Office visit spending	(6) Rx spending	(7) Lab spending
Panel A: omit policy change window (exlude June/July 2008; n = 8840)							
ln(Copay)	−0.196 (0.072)	−0.194 (0.279)	−0.331 (0.161)	−0.231 (0.133)	−0.117 (0.056)	−0.138 (0.082)	−0.300 (0.098)
Panel B: restrict to policy change window (retain April to September 2008; n = 2412))							
ln(Copay)	−0.207 (0.088)	Insufficient data	−0.075 (0.229)	−0.225 (0.140)	−0.251 (0.064)	−0.135 (0.062)	−0.297 (0.091)
Panel C: falsify policy change by changing date (set policy change to January 2008 and keep observations through June 2008; n = 4820)							
ln(Copay)	0.095 (0.096)	0.380 (0.502)	0.072 (0.214)	−0.024 (0.166)	0.129 (0.082)	−0.059 (0.095)	0.047 (0.117)

Note: Each number in the table is an elasticity that comes from a separate GLM regression. Dependent variable is total monthly healthcare expenditure per member per month. Standard errors are clustered on member, and explanatory variables include percent of FPL fixed-effects, and year*month indicators.

this specification is that the loss of data also meant that we were unable to estimate a separate price elasticity of demand for hospital care (there was not enough variation in 6 months of the data to jointly identify all the fixed-effects that we have included and the effect of the copayment change, when it comes to a relatively rare event like hospital care).

Finally, in Panel C of Table 7, we coded the policy change as having occurred on January 1, 2008 (6 months before it actually happened) to verify that we obtained null effects from this placebo regression. Note that the estimated coefficients are never significant and are much smaller in magnitude, thereby reassuring us that we have really picked up the effect of the policy change and not of a pre-existing differential trend.

6. Discussion and implications

In the past, public insurance for our lowest income citizens has featured little patient copayment. This reflects concerns both about affordability and the ability of the poor to effectively manage their care in the face of marginal costs; of particular concerns is the notion that the poor, facing unaffordable copayments, may cut back on necessary as well as unnecessary care. As public insurance subsidies have expanded to citizens above the poverty line, however, fiscal expediency suggests the use of some patient copayments. Yet concerns remain that copayments in this population may do more harm than good.

We have investigated this issue in the Commonwealth Care program in Massachusetts, the nation's largest expansion of publicly subsidized private insurance coverage and the model for the national Affordable Care Act. During our sample period there was a sizeable (in percentage terms) copayment change facing individuals with incomes between one and three times the poverty line. Using a unique set of claims data provided by the state, we are able to study the impacts of this copayment change on overall utilization and on utilization of specific medical services.

Our results largely confirm the conclusions of the RAND Health Insurance Experiment. We find that health care demand is somewhat sensitive to copayments, but that the elasticity is small (−0.16 on average). We find that those who are chronically ill, and especially those with diabetes, high cholesterol and asthma, have a lower price elasticity of demand. We also find no statistically significant evidence for "offset effects" that would indicate that reduced use of outpatient services led to increased demand for hospital services. Among the sickest patients in our sample, there are positive effects on inpatient care, but they are statistically insignificant, so we cannot draw strong conclusions from them. Our results suggest that 'offset effects' in poorer populations are possible, but require more evidence than we are able to provide.

These results are subject to a number of caveats. In particular, these copayment changes are large in percentage terms but small in absolute dollar terms. It is possible that responses could differ for larger copayment changes or changes in coinsurance rates, which would induce larger patient cost-sharing than the copayment changes examined here. Moreover, we have not estimated pure own-price effects here. The decline in prescription drug use, for example, could come from the increase in prescription drug copayments or from the increase in office copayments, which could cause physicians to prescribe fewer drugs. Our aggregate elasticity provides a measure of patient response to an overall rise in copayments, but it does not necessarily speak to the optimal design of insurance for this population. Future work could usefully address these limitations.

An important policy question is the extent to which these results apply to the cost sharing imposed by the Affordable Care Act. For those below 150% of poverty in the ACA, the cost sharing is comparable to that imposed on the Commonwealth Care population studied here. Once income rises above 150% of poverty, however, the cost sharing imposed by the ACA is much larger than that studied here, with enrollees likely facing deductibles of $500 or more as income rises. Future work on the actual impacts of ACA cost sharing changes could usefully corroborate our findings and assess whether they apply to the much higher cost sharing imposed by the ACA on those above 150% of the poverty line.

References

Baicker, K., Goldman, D., 2011. Patient cost-sharing and healthcare spending growth. Journal of Economic Perspectives 25 (2), 47–68.

Baicker, K., Mullainathan S., Schwartzstein, J., 2012. Behavioral Hazard in Health Insurance, NBER working paper number 18468.

Buntin, Melinda Beeuwkes, Zaslavsky, Alan M., 2004. Too much ado about two-part models and transformation? Comparing methods of modeling medicare expenditures. Journal of Health Economics 23 (3), 525–542.

Chandra, A., Gruber, J., McKnight, R., 2010a. Patient cost-sharing and hospitalization offsets in the elderly. American Economic Review 100 (1), 193–213.

Chandra, A., Gruber, J., McKnight, R., 2010b. Patient cost sharing in low income populations. American Economic Review 100 (2), 303–308.

Chandra, A., Gruber, J., McKnight, R., 2011. The importance of the individual mandate: evidence from Massachusetts. New England Journal of Medicine 364 (4), 293–295.

Charlson, M.E., Pompei, P., Ales, K.L., McKenzie, C.R., 1987. A new method of classifying prognostic comorbidity in longitudinal studies: development and validation. Journal of Chronic Disease 40 (5), 373–383.

Congressional Budget Office, 2013. Effects on Health Insurance and the Federal Budget for the Insurance Coverage Provisions in the Affordable Care Act–May 2013 Baseline, Available from www.cbo.gov/publications/44190 on 03.06.13.

Cunningham, P.J., 2002. Prescription drug access: not just a medicare problem. Center for Studying Health System Change 2002 (51), 1–4 (Issue Brief).

Donaldson, C., 2008. Top-ups for cancer drugs: can we kill the zombie for good? British Medical Journal 337, a578, doi:10.1136/bmj.a578.

Evans, R.G., Barer, M.L., Stoddart, G.L., Bhatia, V., 1993. Who are the zombie masters, and what do they want? Vancouver: Centre for Health Services and Policy Research. University of British Columbia.

2013, January 24. Federal Register. U.S. Government Printing Office, Washington, DC.

Gaynor, M., Li, J., Vogt, W.B., 2007. Substitution, spending offsets, and prescription drug benefit design. Forum for Health Economics and Policy 10 (2).

Goldman, D., Joyce, G., Escarce, J., Pace, J., Solomom, M., Laouri, M., Landsman, P., Teutsch, S., 2004. Pharmacy benefits and the use of drugs by the chronically ill. Journal of the American Medical Association 291, 2344–2350.

Goldman, D., Joyce, G., Zheng, Y., 2007. Prescription drug cost sharing: associations with medication and medical utilization and spending and health. Journal of the American Medical Association 298 (1), 61–69.

Goldman, D., Smith, J.P., 2002. Can patient self-management help explain the SES health gradient? Proceedings of the National Academy of Science 99 (16), 10929–10934.

Massachusetts Health Insurance Connector Authority, 2008. Report to the Massachusetts Legislature: Implementation of the Health Care Reform Law 2006–2008, Chapter 58, available from http://www.mahealthconnector.org

Manning, W.G., Newhouse, J.P., Naihua, D., Keeler, E.B., Leibowitz, A., 1987. Health insurance and the demand for medical care: evidence from a randomized experiment. The American Economic Review 77 (3), 251–277.

McGuire, T., 2012. Demand for health insurance. In: Mark, V., Pauly, Thomas, G., McGuire, Pedro, P., Barros (Eds.), The Handbook of Health Economics. Elsevier, pp. 317–396.

Nelson Jr., A.A., Reeder, E.C., Dickson, M.W., 1984. The effect of a medicaid drug copayment program on the utilization and cost of prescription services. Medical Care 22 (8), 724–736.

Newhouse, J., 1993. Free for All: Lessons from the RAND Health Insurance Experiment. Harvard University Press, Cambridge, MA.

Reeder, C.E., Nelson, A.A., 1985. The differential impact of copayment on drug use in a medicaid population. Inquiry 22, 396–403.

Stuart, B., Zacker, C., 1999. Who bears the burden of medicaid drug copayment policies? Health Affairs 18 (2), 201–212.

Tamblyn, R., et al., 2001. Adverse events associated with prescription drug cost-sharing among poor and elderly persons. Journal of the American Medical Association 285 (January (4)), 421–429.

Trivedi, A.N., Moloo, H., Mor, V., 2010. Increased ambulatory care copayments and hospitalizations among the elderly. New England Journal of Medicine 362 (4), 320–328.