Liquidity and Healthcare Consumption: Evidence from Social Security Payments*

Tal Gross[†] Timothy J. Layton[‡] Daniel Prinz[§]

December 2019

Preliminary and Incomplete.
Please do not cite or circulate.

Abstract

Health insurance copayments and deductibles can distort the healthcare consumption of those with no cash. We test for that possibility by studying the distribution of monthly Social Security checks among Medicare recipients. On the day that Medicare recipients receive their Social Security checks, they become 6–12 percent more likely to fill a prescription. That "payday effect" is concentrated in low-income areas, and is of a similar magnitude for a set of drugs where imperfect adherence has important health consequences: blood thinners, insulins, anti-psychotics, and anti-seizure medications. We observe no effect of Social Security checks among Medicare recipients who pay no copays, but we see strong effects for very low-income recipients who face subsidized copayments of only $2–$5. Payday effects disappear when recipients are re-categorized to face no copayments. By contrast, the effects diminish, but do not disappear, when recipients are re-categorized to face lower but non-zero copayments, indicating that only full insurance can eliminate this distortion.

*We are grateful to Rena Conti, Salama Freed, Andrew Goodman-Bacon, Anupam Jena, Michael McWilliams, Joe Newhouse, Analisa Packham, Boris Vabson, and seminar participants at Duke Fuqua, the Federal Reserve Bank of Boston, Harvard Medical School, the University of Hong Kong, the BU-Harvard-MIT Health Economics Seminar, APPAM, ASHEcon, the Notre Dame Health Economics Conference, the Annual Health Economics Conference at the University of California, San Francisco, and the Annual Southeastern Health Economics Study Group at the University of Tennessee, Knoxville for useful feedback. Julia Yates provided superb research assistance. Research reported in this publication was supported by the National Institute on Aging of the National Institutes of Health under Award Number P30AG012810 through the National Bureau of Economic Research Center for Aging and Health Research. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

[†]Boston University and NBER. talgross@bu.edu

[‡]Harvard University and NBER. layton@hcp.med.harvard.edu

[§]Harvard University. dprinz@g.havard.edu

1 Introduction

There exists broad agreement that health insurance ought to involve some form of cost sharing: copayments, coinsurance rates, or deductibles. Theory suggests that cost sharing can limit moral hazard (Zeckhauser, 1970), and empirical evidence confirms that cost sharing reduces spending (Manning et al., 1987; Aron-Dine et al., 2013; Shigeoka, 2014; Brot-Goldberg et al., 2017). Indeed, motivated by that evidence, high-deductible health insurance plans have become increasingly common in the United States. Over only ten years, the share of firms offering such plans increased from 4 percent to nearly 30 percent (Kaiser Family Foundation, 2018). Even some Medicaid programs, which historically have imposed no cost sharing on their recipients, have begun to require small copayments for prescription drugs (Wright et al., 2005).

At the same time, a growing body of research has documented an important pattern regarding the time-path of consumption: the timing of consumption depends on the timing of income. Olafsson and Pagel (2018), for instance, find that low-income consumers purchase 70 percent more goods on days when they receive their paychecks. That finding is consistent with many other studies: consumption responds to both predictable and unanticipated changes in income.[1] One might call such a pattern an "income-timing effect:" consumers, it seems, often delay many of their purchases until their paychecks arrive.

A natural question, then, is how one ought to view a health insurer's cost-sharing mechanisms when its enrollees may exhibit income-timing effects. The optimal insurance contract is typically viewed as one trading off risk protection against moral hazard (Zeckhauser, 1970). However, in the presence of income-timing effects, an additional benefit of generous insurance coverage must be considered: more-generous coverage allows consumers to avoid having to delay their healthcare until the receipt of their income. With generous coverage, recipients can consume healthcare when they need it, rather than waiting for their income to arrive.

There exists, however, little evidence as to whether the timing of income affects the timing of healthcare consumption. To our knowledge, no empirical studies credibly isolate that relationship. And yet, such evidence could lead to substantial implications for policy. If the timing of income does indeed deter or delay welfare-enhancing healthcare consumption, then one might argue that insurers ought to address this

[1] See the work of Parker (1999); Souleles (1999); Agarwal et al. (2007); Agarwal and Qian (2014); Stephens (2003); Johnson et al. (2006); Baugh and Wang (2018); Parker et al. (2013); Gross and Tobacman (2014); Gross et al. (2019).

issue. Insurers could offer financing options that allow consumers to use their future income to cover their current out-of-pocket costs. Alternatively, if the timing of healthcare consumption is unaffected by the timing of income, then such evidence would further strengthen the case for high-deductible plans.

This paper estimates the extent to which the healthcare consumption of Medicare recipients is affected by when they receive their income. We study how the distribution of Social Security checks affects the prescription-drug purchases of Medicare recipients. Since 1997, Social Security recipients have received their monthly checks on a pre-determined schedule based on the day of the month that they were born. This system allows researchers to compare Social Security recipients who just received their check to recipients who are quasi-randomly assigned to receive their checks a week or two later.[2] We compare the prescription-drug purchases of each Medicare recipient around the day they receive their checks to the healthcare consumption of two control groups. First, we study similar Medicare recipients who receive their checks on a different date. Second, we study Medicare recipients who are enrolled in subsidized plans such that they are not charged copayments. These comparisons allow us to isolate the effect of receiving Social Security checks on the consumption of Medicare-covered drugs.

The results, first and foremost, reject the hypothesis that the receipt of income has no effect on healthcare consumption. The probability that a low-income Medicare Part D beneficiary who faces copayments fills a prescription increases by over 10 percent on the day she receives her Social Security check. We observe that effect overall, for all drugs, and also for drugs for which short-term non-adherence can have important health consequences: blood thinners, insulin, anti-psychotics, and anti-seizure medications. In addition, a particular feature of the Social Security check distribution schedule provides an additional source of variation: Two-thirds of checks arrive 28 days after the previous check, and the remaining one-third of checks arrive 35 days after the previous check. We find a larger increase in prescriptions filled when checks arrive after 35 days rather than after 28 days. Finally, we find a clear gradient in the size of the effect by income: Medicare recipients who live in the lowest-income ZIP Codes exhibit income-timing effects that are roughly five times the magnitude of the effects we observe in the highest-income ZIP Codes.

We then assess the extent to which more-generous insurance coverage weakens

[2] Stephens (2003), Evans and Moore (2011), Leary and Wang (2016), and Baugh and Wang (2018) use the distribution of Social Security checks to study outcomes other than healthcare consumption.

these income-timing effects. To do so, we exploit the fact that Medicare recipients often move into more-generous forms of prescription-drug coverage via the Low Income Subsidy (LIS) program. Those recipients transition from facing an average of $19 in copayments per prescription fill to facing $0 in copayments. We find that when recipients undergo that transition, they stop disproportionately filling prescriptions on Social Security paydays. When, by contrast, recipients are re-categorized such that their copayments become smaller but do not shift to zero, we observe a much smaller decrease in the effect of Social Security checks.

We use these transitions onto the LIS program to estimate a price-elasticity of demand for drugs. Our overall estimate is similar in magnitude to what has been documented by previous studies. We then stratify the sample of recipients who transition onto subsidized copayments into two groups: those who previously tended to fill their prescriptions on their Social Security paydays, and those whose prescription fills did not tend to occur on their paydays. We find a much larger price-elasticity of demand for recipients who previously had been filling their scripts on their paydays. This suggests that consumers who exhibit income-timing effects tend to be much more price sensitive.

We consider several different explanations for the increase in prescription-drug fills when Social Security checks are distributed. There are two, related mechanisms that likely drive the findings. First, many Medicare recipients may be liquidity constrained, and so have no choice but to wait until their Social Security checks arrive before they can fill their prescriptions. Indeed, the results are consistent with such a mechanism: we observe a larger effect for more-expensive drugs, and a larger effect for lower-income beneficiaries who are more likely to be liquidity constrained. That said, we also observe effects of Social Security checks even for beneficiaries facing remarkably low copayments of just $2.25 for generic drugs. Thus we cannot rule out the possibility that the effects are also driven by a type of mental accounting. Beneficiaries may not be liquidity constrained overall, but may delay the purchase of healthcare because they have allocated their current liquidity to other purposes.[3]

The final section of this paper explores the implications of these results for the optimal structure of health insurance. First, the results suggest an additional benefit of full insurance. The traditional view is that full insurance reduces the risk that consumers face by reducing the difference in their wealth across states of the world.

[3]An alternative explanation is that recipients wait until they receive their checks to do all of their shopping. We can reject the hypothesis that such a dynamic drives our main results, because we find no effect of Social Security checks for Medicare recipients who do not face copayments.

This paper's results suggest another benefit: full insurance eliminates income-timing effects. Consumers who face no cost sharing do not delay filling their prescriptions until their income arrives. Moreover, the results suggest that this benefit of insurance is non-linear in generosity. Lowering copayments decreases income-timing effects, but the biggest reduction occurs when copayments are eliminated altogether.

Second, the results suggest that income-timing effects may account for part of the overall response of consumers to cost sharing. Under assumptions, discussed below, the results are consistent with a model in which income-timing effects lead consumers to exhibit larger responses to co-payments. We estimate that the price-elasticity of demand is three times larger for Medicare recipients who exhibit income-timing effects than for other, similar recipients. Taken together, these results suggest that the optimal health insurance contract for low-income consumers should be more generous than the typical risk-protection-versus-moral-hazard framework would suggest.

This paper contributes to research that documents the relationship between the timing of income and consumption. Previous studies have demonstrated that consumers often change their behavior after income shocks, even when those income shocks are anticipated. Some studies have exploited the distribution of Social Security checks, as we do. Stephens (2003), for instance, documents that the distribution of Social Security checks affects total consumption, while Leary and Wang (2016) and Baugh and Wang (2018) focus on financial outcomes, such as the takeup of payday loans. Evans and Moore (2011) demonstrate that the distribution of Social Security checks increases short-run mortality. Other studies explore how the arrival of paychecks and government benefits affect food consumption (Shapiro, 2005; Beatty et al., 2019), retailers' responses (Hastings and Washington, 2010; Cheng and Beatty, 2016; Goldin et al., 2019), hospitalizations (Dobkin and Puller, 2007), crime (Foley, 2011), financial decision-making (Carvalho et al., 2016), and aggregate economic activity (Berniell, 2018). Still other studies focus on income shocks that are less anticipated, such as tax rebates and fiscal-stimulus payments. Such studies also uncover significant effects of income receipt on consumption and financial outcomes (Gross and Souleles, 2002; Johnson et al., 2006; Stephens, 2006; Agarwal et al., 2007; Bertrand and Morse, 2009; Agarwal and Qian, 2014; Baugh et al., 2014; Kaplan and Violante, 2014; Carvalho et al., 2016; Baker and Yannelis, 2017; Zhang, 2017; Vellekoop, 2018; Gelman et al., 2019).

The paper is also related to a growing literature on the optimal design of health insurance in particular and social insurance more generally. Most closely related to

this paper, Ericson and Sydnor (2018) explore how health insurance ought to be structured if beneficiaries are liquidity constrained. Baicker et al. (2015) discuss optimal cost-sharing mechanisms when healthcare consumption is shaped by behavioral biases. And, finally, a related issue is the "access motive" for health insurance: the idea that health insurance not only protects against financial risks but also makes otherwise-unaffordable healthcare accessible (Nyman, 1999). These studies grapple with how the traditional economic model of insurance ought to incorporate new findings from the consumption literature and research in behavioral economics. Outside of health insurance, as Chetty and Finkelstein (2013) highlight, most of the previous studies on liquidity and social insurance focus on unemployment insurance (Hansen and Imrohoroglu, 1992; Gruber, 1997; Card et al., 2007; Chetty, 2008; Centeno and Novo, 2009; LaLumia, 2013). This paper shows how these issues should be considered in the context of health insurance and provides empirical evidence that they are important.

The remainder of the paper proceeds as follows. Section 2 provides background on the Social Security benefit schedule, the Medicare Part D program, and the data we use. We then present our empirical results in two sections. First, Section 3 describes how the distribution of Social Security checks affects overall prescription drug purchases. Second, Section 4 describes how Medicare recipients who have their copayments lowered or eliminated change the pattern in which they fill their prescriptions. Finally, Section 5 describes the implications of the results for the optimal structure of health insurance, and Section 6 concludes.

2 Background and Data

2.1 The Distribution of Social Security Checks

Since 1997, Social Security benefits have been distributed to beneficiaries according to a schedule that is based on beneficiaries' dates of birth. Recipients with birthdays on the 1st through the 10th of the month receive their checks on the second Wednesday of each month; recipients with birthdays on the 11th through 20th of the month receive their checks on the third Wednesday of each month; and recipients with birthdays on the 21st through 31st of the month receive their checks on the fourth Wednesday of each month. The only deviations occur when federal holidays fall on Wednesdays. In those rare cases, payments are made one day early.

Two-thirds of checks arrive 28 days after the previous check, but one-third of

checks arrive 35 days after the previous check. Checks would always arrive after 28 days were all months to have precisely 28 days. The irregular structure of the calendar leads to occasional five-week gaps between the second, third, or fourth Wednesday of one month and the corresponding Wednesday of the next month. For example, in 2019 beneficiaries received their paychecks after five weeks in February, June, August, and November.[4]

Since March of 2013, all Social Security recipients have been required to receive their benefits electronically, via direct deposit. Even before that law went into effect, most recipients relied on direct deposit. In early 2012, 94 percent of Social Security checks were distributed via direct deposit (O'Carroll, 2012). The funds thus flow electronically from the Social Security Administration to recipients' banks. Banks then make the funds available to recipients on their scheduled paydays. However, there exists one complication. Some banks provide recipients with the funds the day before their official payday.[5] As a result, below, we find some response to Social Security paychecks the day before payments are supposed to be received.

2.2 Medicare Part D

To study prescription drug consumption, we use data on a 20-percent sample of beneficiaries enrolled in Medicare Part D, a federal program which provides insurance for prescription drugs to elderly and disabled Americans. In Medicare Part D, beneficiaries choose among private plans, but the federal government provides subsidies for premiums and cost sharing depending on beneficiaries' income.

We rely on several datasets for our analysis: the 100-percent Medicare Beneficiary Summary File (BSF), the 20-percent Medicare Part D Drug Event File (PDE), and the Medicare Inpatient (100 percent), Outpatient (100 percent), MedPAR (100 percent), and Carrier (20 percent) claims data. We observe beneficiaries' dates of birth in the BSF which allows us to define the three Social Security payday groups described above. We also observe each beneficiary's geographic location (state, county, and ZIP Code), as well as whether they enroll in Medicare Parts A, B, C, and D. Importantly, we observe, for each month, whether they receive premium and copayment subsidies via the Medicare Part D Low Income Subsidy (LIS) program.

[4]Supplemental Security Income (SSI) payments are made on the first of the month; individuals who are eligible for both Social Security and SSI and individuals who were receiving benefits prior to May 1997 are paid on the third of the month.

[5]We verified this phenomenon in discussions with experts at the Consumer Financial Protection Bureau.

The Medicare Part D Low Income Subsidy program subsidizes premiums and copayments to varying degrees. We divide beneficiaries into three main categories. First, there are beneficiaries who face no copayments. That group consists of individuals enrolled both in Medicare and Medicaid. These individuals receive full premium subsidies and the Medicaid and Medicare programs jointly pay all of their copays. Second, there are beneficiaries who face subsidized copayments. These beneficiaries are sufficiently low income to qualify for subsidized copayments, but not Medicaid. This group actually consists of two separate categories. Some beneficiaries are sufficiently low income to face heavily subsidized copayments: about $2 for generic drugs and about $5 for branded drugs. Other beneficiaries have slightly higher incomes and so receive smaller subsidies. For simplicity, we combine all recipients receiving partial copayment subsidies into one "subsidized-copay" group. Finally, third, there are all other beneficiaries (not enrolled in the LIS program), who face full copayments.[6]

Table 1 presents summary statistics for each of these groups. In the table, we break out the subsidized copay group into its two sub-groups to show the variation in copays within this group, but throughout the paper we combine these groups into one subsidized-copay group. The no-copay group fills slightly more prescriptions than the groups paying copayments. This difference may reflect a demand response to the lower out-of-pocket prices, differences in health status across these groups, other differences in recipients' characteristics, or some combination of these factors. The difference in out-of-pocket prices across the groups paying some copayments is clear in the table, with the average out-of-pocket spending equal to approximately $181 in the subsidized-copay group, and $551 in the full-copay group.

Two other patterns in Table 1 are noteworthy. First, we observe substantially more people in the full-copay group versus the other groups: over 3.7 million people in 2015 in the full-copay group, versus well under a million in the under two groups, combined. That breakdown is representative of the size of these groups in the general population. In the empirical results, below, this leads to much narrower confidence intervals for estimates based on the full-copay group than for estimates based on the other groups.

Second, note that the average copayments faced by beneficiaries in the more highly

[6]There exists an additional group of beneficiaries who pay zero copays and zero premiums. This group consists of Medicaid-Medicare enrolled individuals who are institutionalized. For this group, copays are fully paid by the Medicare program rather than the Medicaid program. We ignore beneficiaries in this group, since they are institutionalized and thus have limited discretion over the timing of their prescription-drug fills.

subsidized component of the subsidized-copay group is remarkably low: an average of $3 per script. That average conceals substantial variation. For this group, 71 percent of fills involve a copayment less than $5.[7] The small magnitude of the out-of-pocket costs these beneficiaries face is important in interpreting some the results below. Even though these beneficiaries face relatively small out-of-pocket costs, some still seem to wait until their payday to fill their prescriptions.

Table 1. Summary Statistics

	Low-income, No copay	Low-income Low copay	Low-income, Higher copay	High-income, Highest copay
Number of beneficiaries in 20% sample in 2015	678,447	102,406	41,373	3,788,347
Mean total scripts per year	52.48	43.23	46.06	29.11
Mean out-of-pocket spending per year	≈$0	$124.49	$320.24	$550.67
Mean share filling a script each day	0.036	0.056	0.064	0.047

Table shows summary statistics for the copay groups used in our analysis. For more details see Section 2.

3 Overall Income-Timing Effects for Prescription Drugs

This section explores how the distribution of monthly Social Security checks affects Medicare-covered prescription-drug purchases. We first examine the total number of scripts filled on Social Security paydays. We then examine how the effect varies for different demographic groups and types of drugs.

3.1 Unadjusted Means

We first study the effect of Social Security checks on prescription drug fills in the raw data. We assign each prescription-drug fill to "event time," that is, to days relative to the nearest date upon which recipients receive their Social Security checks. For each event day, we then calculate the share of recipients filling a script who receive their checks on the second Wednesday of each month. We subtract from that value the corresponding share for recipients who receive their checks on the fourth Wednesday

[7]For the more highly subsidized component of the subsidized-copay group, 66 percent of fills involve a copayment less than $3 and 31 percent of fills involve a copayment less than $2.

of each month. Finally, we calculate the average of that value (across months) for each event day. This calculation has the advantage of differencing out day-of-the-week specific effects.

Figure 1 presents the average differences by event day for two groups: the no-copay group and the subsidized-copay group. The top panel presents the analysis for the share of recipients filling a script on each day. The diamond-shaped markers plot the results of this exercise for the subsidized-copay group. That plot suggests no major differences until the day before Social Security checks are distributed, when the share of beneficiaries filling scripts begins to increase for those approaching payday.[8]

The largest differential increase occurs on the day deposits are made. The share of recipients filling scripts increases by 0.56 percentage points (6.9 percent) on that day relative to other payday groups. The share then returns back to baseline within several days.

As a point of comparison, the figure presents the same exercise for the no-copay group. The circular-shaped markers in Figure 1 plot the re-centered time series for those beneficiaries. For that group, there exists no differential change in the share filling prescriptions when checks are distributed. Importantly, these recipients still receive Social Security checks—the only difference is that they are not required to pay co-payments.

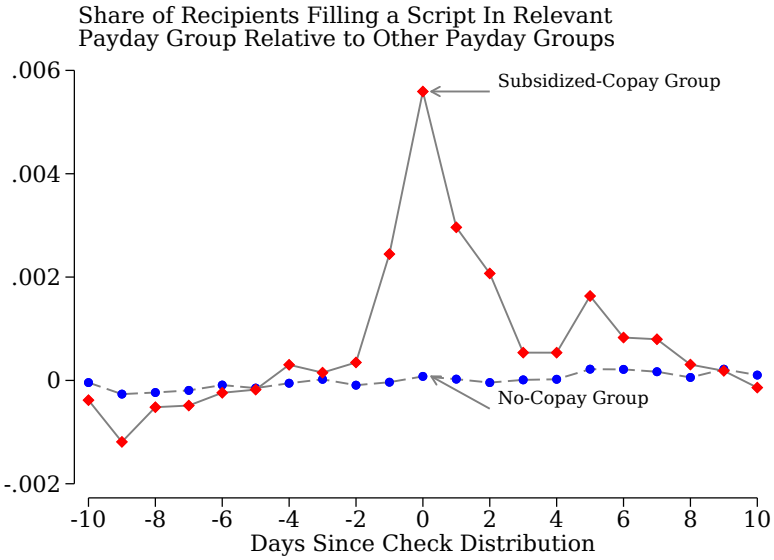
The second panel of Figure 1 repeats this exercise when the outcome of interest is the logarithm of total scripts filled at the group level. We observe the same pattern as in the top panel: an increase the day before checks are distributed, a peak the next day, and then a return to baseline. The numbers suggest that total scripts filled increase by roughly 11.4 percent on the day that checks are sent.

Appendix Figure A1 presents the same analysis for Medicare recipients who face full co-payments. The figure demonstrates the same pattern over time for this group but the magnitude of the increase in filling is roughly half the magnitude for the subsidized-copay group. That difference in magnitude is unsurprising: the full-copay group has much higher income. The top panel shows that the share of recipients filling scripts increases by 0.25 percentage points (3.8 percent) on the day direct deposits are made. The bottom panel shows that total scripts filled increase by roughly 6.2 percent on the same day.

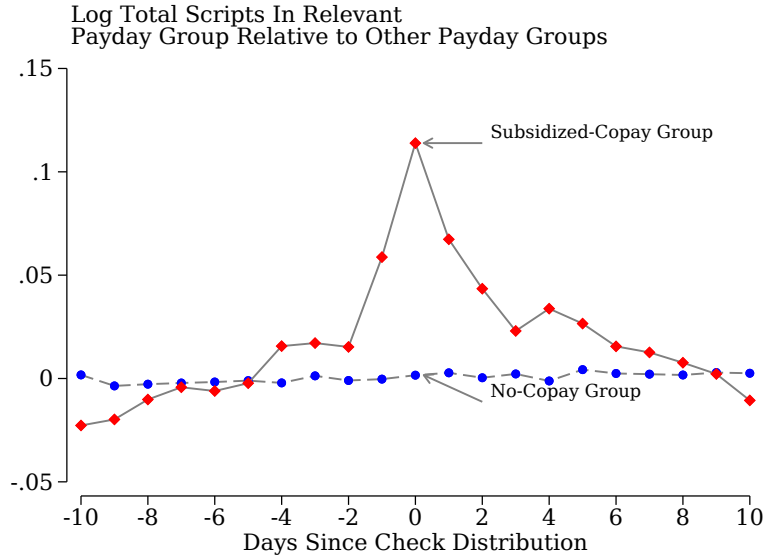
[8]The effects we observe on the day prior to the check disbursement date are consistent with discussions we had with experts at the Consumer Financial Protection Bureau indicating that many banks make these funds available to recipients on the day prior to the actual deposit date.

Figure 1. Re-Centered Time Series

(a) Share of Beneficiaries Who Fill a Script By Day



(b) Logarithm of Total Scripts



Figures plot measures of prescription drug consumption on each event-time day relative to check receipt for the no-copay group and the subsidized-copay group. Panel (a) plots the share of beneficiaries who fill a script. Panel (b) plots log total scripts filled. For each copay group and relative day, we calculate the difference in fills between the beneficiaries who are scheduled to receive benefits versus the beneficiaries who are scheduled to receive their benefits with a two-week difference. For more details, see Section 3.1.

3.2 Regression Approach

We next develop a regression specification, relying on the non-parametric results in Figure 1 to guide us in choosing the appropriate model. Those figures suggest that Social Security checks only affect consumption over a short, three-day period. We thus specify an event-study model that evaluates the effects of Social Security checks within five days of the distribution of the checks.

The regression compares utilization across 31 different birthday groups, with each group defined by the day of month of the Medicare recipient's birthday. Denote the drug consumption on calendar-day t by birthday group $g \in \{1, 2, \ldots, 31\}$ as Y_{tg}. We assign each birthday group an event time for each calendar day, with event time, τ, ranging from -5 to 5, relative to the group's nearest check date. We then estimate the event-study regression:

$$Y_{tg} = \alpha_t + \alpha_g + \sum_{\tau=-5}^{5} \beta_\tau \cdot I_{tg}^{\tau} + \varepsilon_{tg}. \tag{1}$$

Here, I_{tg}^{τ} is an indicator function equal to one if calendar day t is τ days away from the date of Social Security check receipt for individuals in birthday group g.

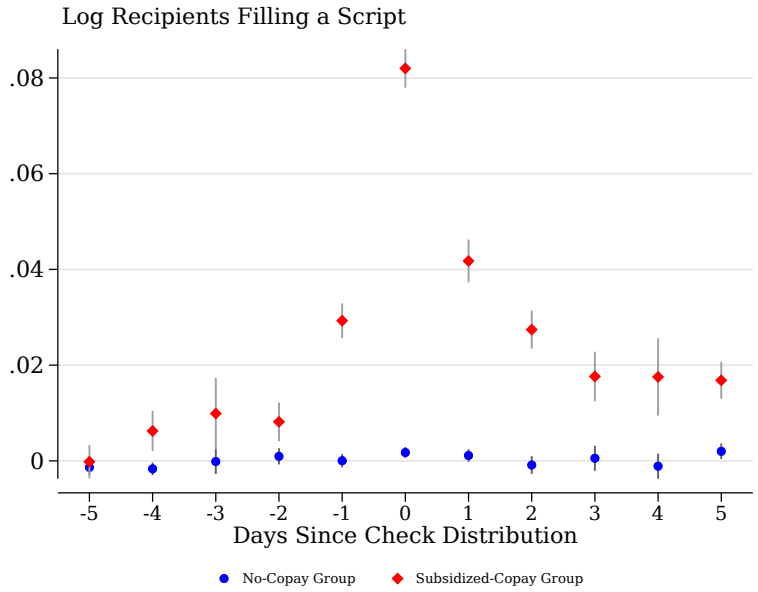
This regression amounts to a standard difference-in-difference specification tracking the consumption of birthday groups over time. To account for auto-correlation, we cluster the standard errors at the level of the birthday group. The objects of interest from this regression are the β_τ's: coefficients that trace out the effect of Social Security checks on behavior just before and just after the checks arrive.

The top panel of Figure 2 presents estimates of the β_τ's for the subsidized-copayment group when the outcome of interest is the log number of recipients filling a script.[9] The results are qualitatively and quantitatively similar to our findings in Figure 1. The point estimates suggest a statistically significant effect of Social Security checks. The number of beneficiaries filling a script begins increasing on day -1, hits a maximum on payday, and then declines. The magnitude of the day-zero estimate, β_0, is 8.2 percent. The bottom panel of Figure 2 presents the same exercise when the logarithm of total scripts is the outcome of interest. The magnitude of the day-zero estimate, β_0, suggests a differential increase of 11.3 percent over the baseline level of daily prescription fills.

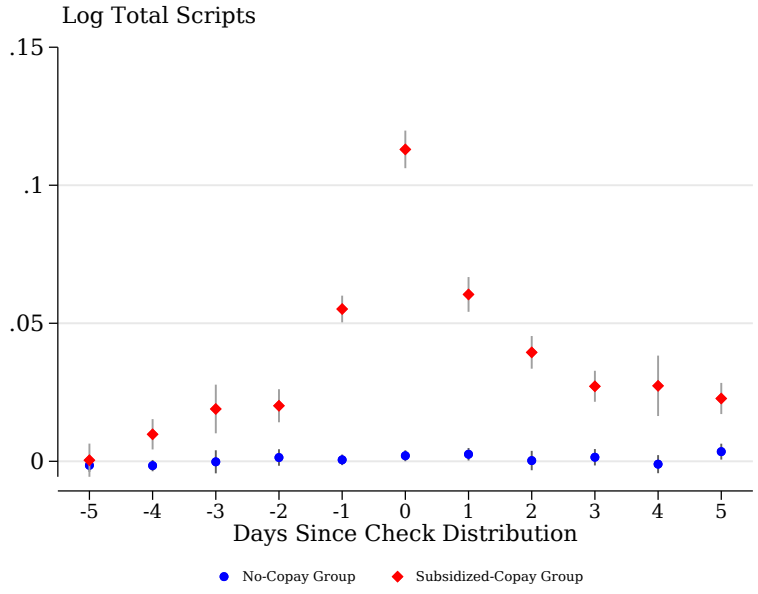
[9] Appendix Figure A2 presents the same event study for the full-copayment group.

Figure 2. Event Study

(a) Log Beneficiaries Who Fill a Script By Day



(b) Logarithm of Total Scripts



Figures plot event study estimates of changes in measures of prescription drug consumption on each event-time day relative to check receipt for the no-copay group and the subsidized-copay group. Panel (a) plots estimates for the share of beneficiaries who fill a script. Panel (b) plots estimates for log total scripts filled. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details, see Section 3.2.

Figure 3 plots estimates of β_0 from Equation (1) when the logarithm of total fills is the outcome of interest. The figure plots those estimates for each of three groups: those who face no copayments, those who face subsidized copayments, and those who face full copayments. For the group that does not pay copayments, we estimate a precise zero effect; the confidence interval rules out increases in fills larger than 0.4 percent. The effects for the other two groups are both economically and statistically significant: we estimate a "payday effect" of 11.3 percent for the subsidized-copay group and an effect of 6.1 percent for the full-copay group.

Figure 3 also presents effects across the two different kinds of Social Security payments: those that arrive after a 28-day wait and those that arrive after a 35-day wait. We observe a larger increase in prescription-drug fills when it has been longer since beneficiaries' last check was sent. For the subsidized-copay group, for instance, we observe an effect of 10.6 percent for checks arriving after a 28-day wait vs. an effect of 12.4 percent for checks arriving after a 35-day wait.

Figure 3. Summary Figure

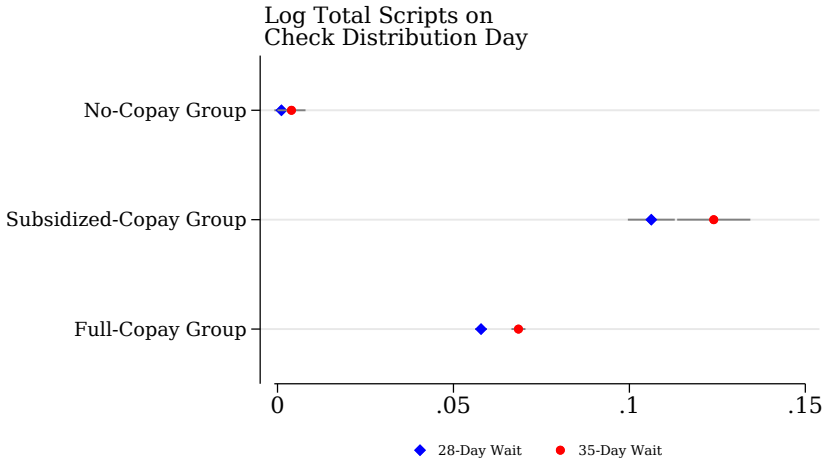


Figure shows the estimated event study coefficient on check distribution day for log total scripts filled by copay group (no-copay group, subsidized-copay group, or full-copay group) and by length of wait (28-day wait or 35-day wait). Standard errors are clustered at the level of the birthday group. These coefficients are from estimating Equation (1). For more details, see Section 3.2.

3.3 Heterogeneous Effects

A remaining question involves treatment-effect heterogeneity. How does the effect of Social Security checks differ by the importance or expense of the drug in question? Are effects larger for certain groups of recipients? To address these questions, we estimate Equation (1) separately for different types of drugs.

Figure 4 plots the percent-increase in scripts when checks are distributed across different types of drugs. The first two rows suggest that, on Social Security paydays, fills increase more for branded drugs than for generic drugs. Next, we stratify drugs by whether or not they involve a large out-of-pocket cost for recipients.[10] The results suggest a larger effect of Social Security checks for expensive drugs. Since branded drugs are typically more expensive, these two patterns are consistent with one another.

Figure 4. Heterogeneous Effects

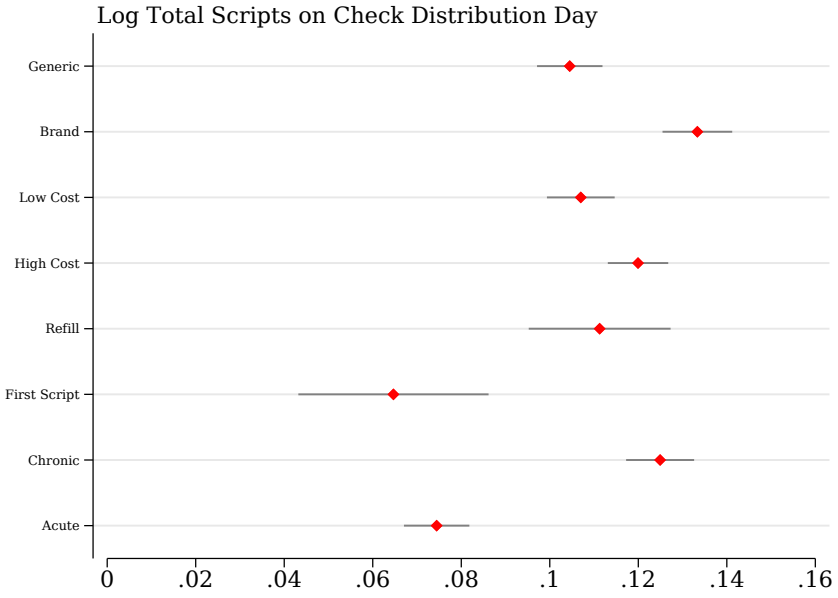


Figure shows the estimated event study coefficient on check distribution day for log total scripts filled by drug type (generic or brand), drug cost (above or below median cost), type of fill (first fill or refill), and frequency of fill (chronic or acute) for the subsidized-copay group. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details, see Section 3.2.

Next, Figure 4 presents estimates when we stratify prescription drug purchases by whether the purchase is a recipient's first fill of the drug, or whether it is a refill

[10]For each drug, we calculate the average out-of-pocket cost for all Medicare recipients. We then categorize a drug as expensive if the average out-of-pocket cost is above the mean of this variable.

of their prescription. These results indicate that there is a significant payday effect for both refills and initial fills. While recipients may have previously stockpiled pills for drugs they are re-filling, this is unlikely to be the case for an initial fill of a new drug. The payday effect for initial fills thus provides clear evidence that recipients are going without drugs for at least some period of time.

We also assess whether different types of recipients are more likely to drive the increase in fills on Social Security paydays. Figure 5 divides the sample by income. The figure plots estimates of β_0 for each ventile of ZIP Code income for recipients in the full-copay group.[11] The figure suggests a clear gradient, with larger effects of Social Security checks for lower-income recipients. Indeed, the payday effect for the lowest income ventile (8.3 percent) is more than four times as large as the payday effect for the highest income ventile (1.8 percent).

Figure 5. Heterogeneous Effects By ZIP Code Income

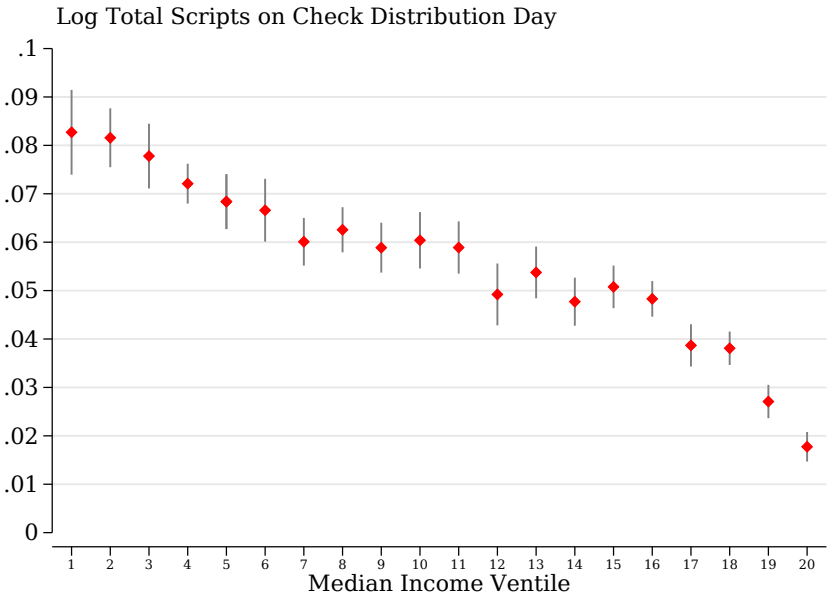


Figure shows the estimated event study coefficient on check distribution day for log total scripts filled by ventile of median zipcode income for the full-copay group. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details, see Section 3.2.

Finally, Figure 6 examines the effect of Social Security checks for particular drugs. We interviewed a panel of physicians and asked for examples of common drugs that

[11] We link every ZIP Code of residence in the data to the median household income as measured by American Community Survey, 2006-2010.

lead to severe, short-term consequences if patients do not adhere to their prescriptions. The physicians suggested the following drugs: blood thinners – used to prevent stroke among patients at high risk for blood clots; insulin – used to regulate blood sugar among diabetics; anti-psychotics – used to treat psychotic disorders such as bipolar disorder and schizophrenia; and anti-seizure medications – used to limit seizures among patients at high risk for seizures. For all of those types of drugs, Figure 6 suggests sizable effects.

Figure 6. Heterogeneous Effects By Drug

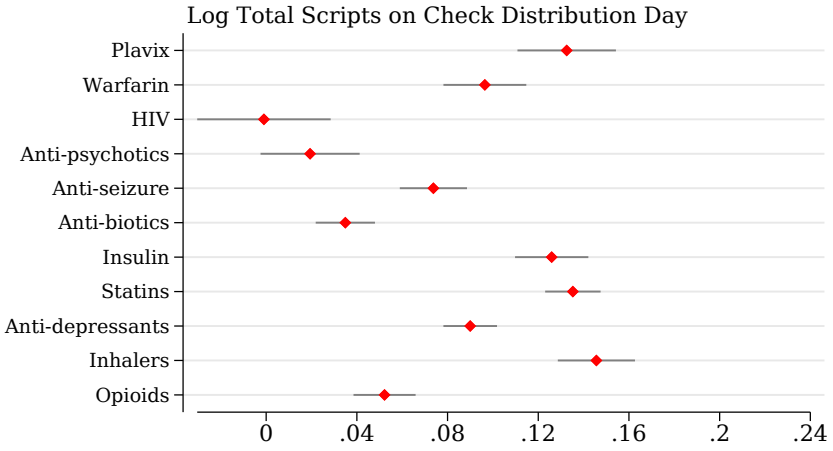


Figure shows the estimated event study coefficient on check distribution day for log total scripts filled of specific drugs and drug classes for the subsidized-copay group. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details, see Section 3.2.

It is important to recognize that interpreting Figures 3, 4, 5, and 6 is a non-trivial task. On the one hand, the figures exhibit patterns that are consistent with liquidity constraints. Low-income consumers are more likely to become liquidity constrained, and expensive drugs are more likely to be difficult for liquidity-constrained consumers to purchase. Indeed, the figures above suggest a larger effect for those types of consumers and those types of drugs. At the same time, the estimates summarize treatment effects based on different populations, as the groups taking different drugs differ from each other in observed and unobserved ways. Thus, the estimated effects are not directly comparable.

4 The Effect of Reduced Copayments on Income-Timing Effects

The results above suggest an increase in prescription drug fills when consumers receive their Social Security checks, ranging from 6.2 percent for (higher-income) recipients who face full copayments to 11.4 percent for (lower-income) recipients who face subsidized copayments. This section studies what occurs when such consumers experience a reduction in the copayments they face. That reduction in copayments stems from transitions from the full-copayment group to either the subsidized-copayment group or the no-copayment group.

We study those transitions to more-generous coverage as a way to assess how income-timing effects relate to the generosity of health insurance. More-generous coverage may allow some beneficiaries to consume healthcare when they need it rather than waiting for their paychecks to arrive. In that case, such a mechanism would amount to an additional benefit of generous coverage, beyond the standard risk-protection benefits of health insurance.

We test for such a dynamic by leveraging within-person transitions across levels of coverage provided by the Part D LIS program. As detailed in Section 2, above, the LIS program consists of several tiers of benefits, with the benefit tier determined by the beneficiary's income and assets.[12] As above, we group the levels of coverage into three tiers: no-copay, subsidized-copay, and full-copay. Many Medicare Part D beneficiaries enter Part D not enrolled in the LIS program but enroll in the program at a later date. We do not observe the reason why a beneficiary enrolls in LIS. Their enrollment likely occurs after either (1) they learn about eligibility for the program or (2) they experience a change in their income or assets that makes them eligible for LIS.

4.1 The Effect of Lower Copayments on Payday Filling

We start by identifying all beneficiaries who switch to the no-copay group from either the subsidized-copay group or from the full-copay group between 2007 and 2015. We

[12] Eligibility for most tiers is partially linked to eligibility for full and partial Medicaid benefits. Full Medicaid beneficiaries (full-benefit dual eligibles) are in the no-copay group. Partial-benefit duals (Qualified Medicare Beneficiaries [QMBs], Specified Low-Income Medicare Beneficiaries [SLMBs], and Qualifying Individuals [QIs]) are each eligible for different non-zero, but highly subsidized tiers of the LIS program. All of these individuals are automatically enrolled in the LIS program when they are enrolled in the other low-income-support program. Additionally, some individuals not eligible for these other programs are eligible for the LIS program.

restrict to beneficiaries continuously enrolled in Part D for at least 1 year before the transition and at least 1 year after. That restriction allows us to observe pre- and post-transition consumption patterns. Appendix Figure A3 plots the time series of enrollment in each of the three groups. The four panels of the figure suggest that membership in a particular group is somewhat persistent. Three years after a transition to the no-copay group, around 80 percent of the switchers remain in the no-copay group, with 10 percent in the full-copay group and 7 percent in the subsidized-copay group (around 3 percent drop out of Part D altogether). Additionally, a year prior to the transition, about 60 percent of switchers were in the full-copay group, while about 40 percent were in the subsidized-copay group. Because the pre-transition copay levels are so different for beneficiaries coming from each of these two groups, in some of our analyses, below, we analyze them separately.

We then determine whether the income-timing effects documented by Section 3 decrease when a beneficiary switches to the no-copay group, that is, when they move from partial insurance to full insurance. To do so, we calculate the share of Wednesday fills that occur on their payday. We focus on the share of *Wednesday* fills that occur on payday, because there are day-of-the-week effects in prescription filling behavior. We calculate the share of Wednesday fills occurring on paydays for each month, across all beneficiaries who switch groups.

If recipients fill prescriptions on a schedule unrelated to their Social Security paydays, then roughly 0.23 of their Wednesday fills should occur on their Social Security payday. That number stems from a simple calculation: there are 12 Social Security checks sent each year, and 52 or 53 Wednesdays in the year. If, however, some consumers delay filling scripts until their paychecks arrive, then the share of their Wednesday fills that occur on their payday would rise far above 0.23.

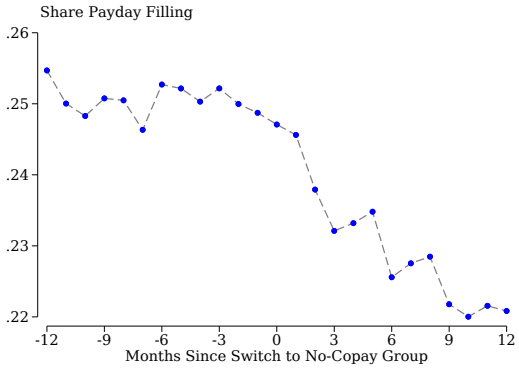
The first panel of Figure 7 plots this measure of the share of Wednesday fills that occurred on payday for each month relative to the month that recipients transition onto the no-copay group. Prior to the transition into the no-copay group, recipients fill roughly 0.25 of their Wednesday fills on their Social Security payday. In the months immediately following their transition to full insurance, however, the share drops to roughly 0.225. Full insurance, it seems, reduces the propensity for recipients to fill their scripts on their Social Security payday.

The remaining panels of Figure 7 perform the same exercise but separately for those shifting from the subsidized-copay group (Panel B) and those shifting from the full-copay group (Panel C). Both panels suggest the same basic pattern: a drop

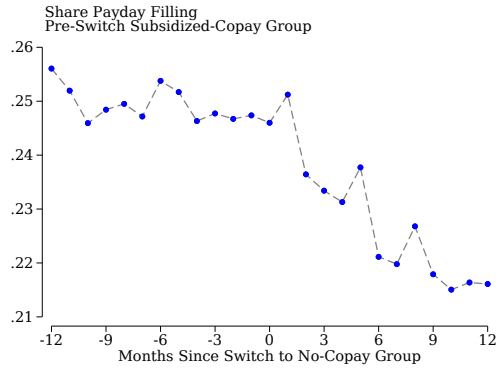
in the share of fills on payday once the recipient is moved into the no-copay group. Further, the effect seems similar for the recipients coming from the full-copay group and recipients coming from the subsidized-copay group.

Figure 7. Share Filling Scripts on Payday

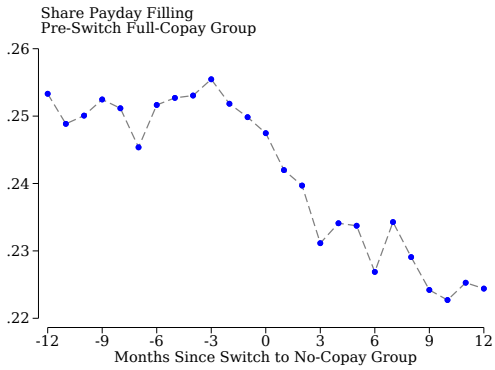
(a) Switch to the No-Copay Group



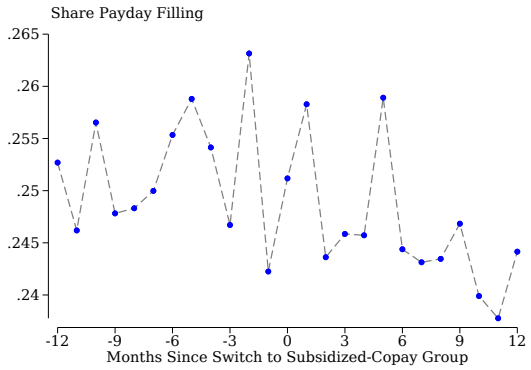
(b) Switch From Subsidized-Copay Group to No-Copay Group



(c) Switch From Full-Copay Group to No-Copay Group



(d) Switch From Full-Copay Group to Subsidized-Copay Group



Figures show the share of Wednesday fills that fall on payday by month relative to switching to the no-copay group. Panel (a) shows results for beneficiaries switching to the no-copay group from either the subsidized-copay group or the full-copay group. Panel (b) shows results for beneficiaries switching to the no-copay group from the subsidized-copay group. Panel (c) shows results for beneficiaries switching to the no-copay group from the full-copay group. Panel (d) shows results for beneficiaries switching to the subsidized-copay group from the full-copay group. For more details, see Section 4.1.

Finally, the last panel of Figure 7 studies recipients switching from the full-copay group to the subsidized-copay group. Recall, that switching between these groups leads to a large reduction in average copayments (from \$19 to \$3 on average), but it does not fully eliminate copayments, unlike transitions into the no-copay group. The figure suggests a pre-switch rate of payday filling of roughly 0.25, similar to those switching from the full-copay group to the no-copay group. However, we observe a small, gradual decrease in payday filling, settling to 0.24. That change is substantially smaller than what we observe for recipients transitioning to full insurance.

In order to compare these changes in payday filling, we estimate event-study and difference-in-difference regressions. We assign beneficiaries to cohorts based on the month that they transition onto the no-copay group. We also construct a control group of beneficiaries who never switch copay groups, in order to remove calendar-month-specific effects. Because our measure of payday filling is an aggregate and not an individual measure, we collapse the data to the cohort-level, and consider the control group as one large non-switching cohort. We then run the following regression at the cohort-month level:

$$\text{Payday Filling}_{ct} = \sum_{\tau=-36}^{36} \beta_\tau \times \mathbf{1}[t - s_c = \tau] \times \text{Switcher}_c + \gamma_c + \delta_t + \varepsilon_{ct}. \qquad (2)$$

Here, $\text{Payday Filling}_{ct}$ is the portion of Wednesday fills that occurred on payday for cohort c in calendar month t. The variable s_c is the calendar month that recipients in cohort c transitioned into the no-copay group; Switcher_c is an indicator function equal to one if the cohort had switched groups; γ_c is a cohort-specific fixed effect; and δ_t is a calendar-month-specific fixed effect. The β_τ's are the coefficients of interest, estimating how the difference in the outcome changes over time relative to the switch date. We also estimate a simplified version of Equation (2) in which we substitute the event-study indicator functions for a single indicator function equal to one if the cohort has already switched to lower copayments.

Column 1 of Table 2 and Appendix Figure A4 present estimates of Equation (2). The results are consistent with the pattern demonstrated by Figure 7. When recipients transition from the full-copay group to the no-copay group, payday filling decreases by 2.1 percentage points (8.5 percent).

Table 2. Switcher Regression Table

	(1) Payday Filling	(2) Portion Out-of -Pocket	(3) Out-of -Pocket	(4) Log Out-of -Pocket	(5) Total Out-of -Pocket	(6) Log Total Out-of -Pocket	(7) Total Scripts	(8) Log Total Scripts	(9) Total Spend	(10) Log Total Spend
Post-Switch	-0.0211***	-0.1935***	-5.1691***	-0.4267***	-17.7503***	-0.4034***	0.4437***	0.1153***	56.4964***	0.3775***
	(0.0022)	(0.0009)	(0.0393)	(0.0028)	(0.1466)	(0.0039)	(0.0065)	(0.0017)	(1.5125)	(0.0051)
Beneficiary FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year × Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cohort FE	Yes	No	No	No	No	No	No	No	No	No
Mean	0.2481	0.2982	6.3883	1.0764	22.4873	1.5831	2.2308	0.8323	119.5361	2.6305
N	12,085	6,670,584	10,031,436	10,031,436	10,031,436	10,031,436	10,031,436	10,031,436	10,031,436	10,031,436

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table shows estimates of the impact of switching to the no-copay group on a variety of outcomes, including payday filling (Column (1)), the portion of costs that is out-of-pocket (Column (2)), out-of-pocket spending (Column (3)), total scripts filled (Column (4)), log total scripts filled (Column (5)), total spending (Column (6)), and log total spending (Column (7)). These coefficients are from estimating Equation (2). For more details see Section 4.1.

4.2 The Effect of Lower Copayments on Demand

We next use an individual-level version of Equation (2) to study the effect of lower copayments on other outcomes: total prescriptions filled, the portion paid out-of-pocket, the total out-of-pocket payment, and total spending. Columns 2–10 of Table 2 present estimates of the individual-level difference-in-difference regression with those variables as outcomes.[13] The table suggests that the transition to lower copayments increases prescriptions filled and total spending while the out-of-pocket payment decreases. These estimates imply an arc elasticity of demand for prescription fills of -0.25.[14] This elasticity is consistent with prior research on the demand for prescription drugs (Chandra et al., 2010; Einav et al., 2015).

As shown in Section 4.1, full insurance relaxes the constraint that led to recipients timing their prescription fills to check receipt. Thus, full insurance may lead to an increase in drug fills both due to moral hazard and due to the relaxation of the constraint. To explore this possibility, we next estimate how elasticities differ for beneficiaries who tend to fill their prescriptions on their Social Security paydays. In order to do so, we calculate, for each recipient, the proportion of their Wednesday fills that occur on their Social Security payday during the period prior to their transition to the no-copay group.[15] We then group recipients by deciles of this measure, the share of their Wednesday fills that occur on their payday.[16]

Table 3 presents estimated elasticities for those groups of recipients. The elasticities we estimate for recipients in deciles 1-7 are (mostly) in the typical range of -0.1 to -0.2. However, for the "payday fillers" or those in deciles 8-10, the elasticities are much larger, around -0.3 for deciles 8 and 9 and -0.6 for the top decile. Thus, the recipients most likely to fill prescriptions on payday prior to switching exhibit a price

[13] Appendix Figure A4 presents event study versions and Appendix Figure A5 presents re-centered time series for those outcomes, as well.

[14] We calculate the arc elasticity as the percent change in the number of prescriptions filled in a month divided by the percent change in the out-of-pocket payment, as in Chandra et al. (2010).

[15] For this exercise, we eliminate from the sample recipients who have fewer than 20 percent of their Wednesday fills occurring on their payday. We do so, because those recipients likely are timing their fills against some other source of income. For instance, they may be relying on their spouse's Social Security check as the relevant "check date" for the household, they may be receiving regular checks through the Supplemental Security Income program, or they may have a regular paycheck from an employer.

[16] To understand whether the share of Wednesday fills that occur on a payday is a plausible measure of the extent to which a recipient times her drug fills to her income, in Appendix Figure A6 we plot our main event study estimates for the subsidized-copay group by quartile of the share of Wednesday fills that occur on a payday. We estimate this share as a leave-one-out measure, calculating it over all years except the year for which we estimate the event study. This figure suggests that beneficiaries who fill a high proportion of their scripts on payday exhibit larger income timing effects.

Table 3. Demand Elasticity By Proportion of Wednesday Fills on Payday

	(1) Δ Total Scripts	(2) Δ Average OOP	(3) Elasticity
All	0.44370***	-5.16913***	-0.24581
	(0.00653)	(0.03926)	
Proportion of Wednesday Fills on Payday > 0.2			
Decile 1	0.41221***	-6.12504***	-0.18887
	(0.03295)	(0.17781)	
Decile 2	0.30036***	-6.74003***	-0.12538
	(0.02571)	(0.17481)	
Decile 3	0.18708***	-6.95178***	-0.06565
	(0.05687)	(0.29257)	
Decile 4	0.27312***	-7.00561***	-0.10725
	(0.03166)	(0.18950)	
Decile 5	0.38229***	-6.07878***	-0.18531
	(0.02641)	(0.17107)	
Decile 6	0.29173***	-7.22395***	-0.11512
	(0.03933)	(0.28365)	
Decile 7	0.37557***	-6.85654***	-0.16292
	(0.03299)	(0.21995)	
Decile 8	0.52801***	-5.68107***	-0.30259
	(0.02587)	(0.16942)	
Decile 9	0.53694***	-5.88397***	-0.27822
	(0.03189)	(0.18100)	
Decile 10	0.71103***	-3.45659***	-0.61207
	(0.02927)	(0.14831)	
Proportion of Wednesday Fills on Payday < 0.2 or Missing	0.49110***	-4.28414***	-0.32694
	(0.00867)	(0.04944)	

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table shows estimated demand elasticities by the proportion of Wednesday fills that fall on a payday. The first row shows results for all beneficiaries. The subsequent rows limit to beneficiaries who have at least one-fifth of their Wednesday fills on a payday and breaks them up into deciles of the proportion of Wednesday fills on payday. The final row shows results for beneficiaries who have less than one-fifth of their Wednesday fills on a payday. Column (1) shows the change in total scripts filled when switching to lower co-payments. Column (2) shows the change in average out-of-pocket spending when switching to lower co-payments. Column (3) shows the implied demand elasticity. These coefficients are from estimating an individual-level version of Equation (2). Standard errors are clustered at the level of the individual. For more details, see Section 4.2.

elasticity more than two times the elasticity of the average recipient. As we discuss in Section 5.3, this pattern is consistent with, though not conclusive evidence of, some part of any demand response to a price change among low-income recipients being due to the relaxation of a constraint rather than moral hazard.

5 Mechanisms and Implications for Health Insurance

The evidence above suggests that low-income Medicare beneficiaries are more likely to fill prescriptions on their Social Security paydays. When beneficiaries move from less-generous insurance to full insurance, that "payday effect" is largely eliminated. Finally, beneficiaries' "demand response" of moving from less- to more-generous insurance is larger for groups that exhibit high rates of payday filling prior to the switch to full insurance.

This section assesses the implications of those results for the economics of health insurance. First, we consider potential mechanisms behind the results. Second, we discuss the implications of the results for the value of insurance. Third, we discuss the implications of the results for the conventional view of moral hazard.

5.1 Potential Mechanisms

The results imply that, for some beneficiaries, it is more costly to fill prescriptions at the end of a pay period than at the beginning of a pay period. We posit three potential explanations for that pattern: increasing returns to shopping, liquidity constraints, and mental accounting.

A first potential explanation is that consumers tend to do all of their shopping at once for the sake of returns to scale. For instance, suppose that some consumers have to pay a bus fare in order to travel to the pharmacy. In that case, they may choose to visit the pharmacy at the same time as they do their other shopping. And thus, they would delay filling their prescriptions until they were prepared to purchase many other goods, as well.

Such an explanation, however, is easy to reject. We find no effect of Social Security checks for the no-copay group. Those consumers still receive Social Security checks, and they, presumably, still face the same potential returns to scale from shopping – they still have to pay a fare for the bus. Nevertheless, that group does not exhibit income-timing effects for prescription drugs, and beneficiaries who transition to the group stop exhibiting income-timing effects once the copayments they face are

eliminated.

A second explanation is that some Medicare beneficiaries are liquidity constrained. At the end of each pay period, some beneficiaries have little cash on hand. Those consumers may face a steep price to access additional liquidity. They may, for instance, face high interest rates from lenders, or have already fully utilized their credit cards. Such consumers would then have no choice but to wait for their next Social Security check before they could fill their prescriptions.

On the one hand, the results are consistent with such an explanation. The consumption literature suggests that lower-income consumers are more likely to be liquidity constrained (Olafsson and Pagel, 2018), and indeed, Figure 5, above, suggests that the effects of Social Security checks documented above are driven by lower-income Medicare recipients. Moreover, more-expensive purchases are more likely to be unattainable for liquidity-constrained consumers, and Figure 4, above, suggests that the effects are larger for more-expensive drugs.

That said, it is surprising that the small copayments faced by recipients in the subsidized-copay group would pose such an obstacle. Thus, we cannot reject a third, alternative explanation: the income-timing effects we observe may be driven by behavioral biases. In particular, Medicare beneficiaries may delay filling their prescriptions based on a type of mental accounting (Thaler, 1985, 1999). They may not be liquidity constrained overall, so much as that the mental account that they devote to medical expenses may be empty.

Unlike liquidity constraints, mental accounting does not imply large *financial* costs to accessing the cash necessary for copayments; rather, it implies large *psychological* costs. The results are consistent with such an explanation. In particular, consider the finding that payday filling is nearly eliminated when recipients transition from the subsidized-copay group to the no-copay group while there is only a small effect from transitioning from the full-copay group to the subsidized-copay group. In that sense, making the product free to the consumer has a much larger effect on the timing of prescription fills to income than making the product cheaper. Such a pattern is consistent with extreme liquidity constraints that prevent purchases even when the price is only $2, but it is also consistent with mental accounting.

In summary, we cannot reject either liquidity constraints or mental accounting as the mechanism driving the income-timing effects we document. It is possible that both mechanisms are empirically relevant, if only for different groups of recipients.

Importantly, though, the two explanations lead to different implications for policy.

If liquidity constraints are responsible for the inter-temporal distortions we observe, then providing beneficiaries with access to short-term credit against their future Social Security checks may be beneficial. If mental accounting is the explanation, however, such access to credit will have little effect on payday filling. Instead, moving from monthly Social Security payments to smaller weekly payments may be an appropriate policy response. By setting payments such that beneficiaries are more likely to have cash-on-hand at any given point in time, policymakers can help beneficiaries avoid the psychological constraints related to the end of the pay period and fill their prescriptions on time.[17]

5.2 The Consumption-Smoothing Benefits of Insurance

Regardless of whether liquidity constraints or mental accounting drive the income-timing effects we document, the results suggest a benefit of insurance that is often not discussed. Typically, insurance is viewed as providing risk protection at the expense of moral hazard (Zeckhauser, 1970). As Ericson and Sydnor (2018) have argued, generous insurance provides an additional benefit: it allows consumers to avoid income-timing effects.[18]

The results above provide direct evidence for that benefit. Low-income Medicare recipients exhibit income-timing effects, but those that face no copayments do not. When recipients are re-categorized such that the copayments they face decrease to zero, they cease to exhibit such effects. In other words, when consumers are provided with full insurance, their healthcare consumption ceases to depend on the timing of their income.

Further, this setting allows us to compare the effects of lowering copayments substantially versus lowering copayments to zero. That comparison, in Table 2, above, suggests that these benefits of insurance are non-linear in the generosity of coverage. The results suggest that much of the consumption smoothing benefit of insurance is concentrated at the transition from almost full insurance to full insurance.

These results suggest the need for a new framework for evaluating policies related to health insurance and cost sharing. The conventional framework evaluates policies as trading off the costs of moral hazard relative to the benefits of risk protection.

[17] Baicker et al. (2015) discuss optimal cost-sharing mechanisms when healthcare consumption is shaped by behavioral biases.

[18] A related issue is the "access motive" for health insurance: the idea that health insurance not only protects against financial risks but also makes otherwise-unaffordable healthcare accessible (Nyman, 1999).

This paper's results suggest that consumption-smoothing benefits must also be acknowledged as an additional benefit of generous coverage, especially in low-income populations. That benefit appears to result mostly from full insurance, rather than being linear in the generosity of coverage.

5.3 The Conventional Interpretation of Demand-Response as Moral Hazard

The theoretical (Arrow, 1963; Pauly, 1968; Zeckhauser, 1970) and empirical (Manning et al., 1987; Finkelstein et al., 2012) literature on health insurance typically interprets a demand response to changes in cost sharing as evidence of moral hazard (Einav and Finkelstein, 2018). That interpretation, however, depends on the assumption that individuals can inter-temporally reallocate income. If, instead, individuals face financial or psychological constraints that prevent them from shifting money over time, then a demand-response to more-generous insurance cannot be interpreted as solely a moral-hazard effect.[19] Instead, the demand-response to more-generous insurance consists of both a moral-hazard effect and what might be called a "timing effect." Both effects imply that individuals consume more healthcare as the out-of-pocket price decreases. When the out-of-pocket price decreases, the moral-hazard effect leads individuals to consume some healthcare that they value less than its marginal cost because they pay less than its marginal cost. At the same time, when the out-of-pocket price decreases, the timing effect means that individuals consume some healthcare they value at least as much as its marginal cost but they otherwise could not allocate the liquidity to pay for it.

This distinction between moral-hazard effects and timing effects is important because the moral-hazard component of the demand response is welfare decreasing. Moral hazard imposes a social cost that is greater than the private cost. Meanwhile, the timing component of the demand response is welfare enhancing because it corrects a credit-market failure or behavioral bias. In other words, when exhibiting moral hazard, individuals "over-consume" or consume healthcare that they value below the social cost. When exhibiting a timing response, individuals do not "over-consume" but instead increase consumption of healthcare for which they would have an *unconstrained* willingness-to-pay that is higher than its marginal cost.

Our results in Table 3 indicate that the individuals who are most likely to time their drug consumption to their income receipt (i.e. the "payday fillers" in the top

[19] Chetty (2008) makes this point in the conext of unemployment insurance.

decile) also have the highest demand elasticities. Indeed, their elasticities are more than two times the elasticity of the average recipient switching to the no-copay group. Further, if we make the (very strong) assumption that in the absence of the liquidity/psychological constraints that cause deciles 8-10 to time their drug fills to income receipt these recipients would exhibit price elasticities equal to the average elasticity among the recipients in deciles 1-7, the overall elasticity would shrink by 37%. If we make the weaker assumption that the "unconstrained" elasticity among deciles 8-10 is *at least as large* as the average elasticity we estimate for deciles 1-7, then 37% would represent an upper-bound on the portion of demand-response that is due to the relaxation of these constraints.

In summary, our results provide suggestive (but not conclusive) evidence that some portion of estimated demand-response to a drug price change among low-income recipients is likely to be the efficient consequence of relaxing a liquidity or psychological constraint, rather than moral hazard. It is important to note, however, that our results do not conclusively prove that this is the case: It could be the case that payday fillers have very different preferences than non-payday fillers, placing much lower valuation on prescription drugs. This type of preference heterogeneity would result in similar heterogeneity in elasticities, and we clearly cannot rule out this possibility. We do believe, however, that these results raise important questions about the conventional interpretation of demand-response as inefficient moral hazard.

6 Conclusion

The results above suggest that Medicare Part D recipients delay the purchase of prescription drugs until their Social Security checks arrive. We estimate a 10.6-12.4 percent increase in prescription fills among low-income recipients facing subsidized copayments on the day checks arrive. We also estimate a 6.1 percent increase among high-income recipients facing full copays. This increase is driven by recipients living in low-income zip codes. It is stronger for more-expensive drugs and branded drugs. There is no payday effect for low-income recipients whose copays are subsidized, providing strong evidence that the effects we observe are directly connected to liquidity-related delays in drug consumption, not spillovers related to other forms of consumption.

The results also suggest a re-interpretation of some classic results in health economics. It has long been recognized that consumers' demand response to the price

of healthcare may not indicate moral hazard, but rather may suggest liquidity constraints bind. Empirical evidence, however, of the existence and extent of liquidity constraints in healthcare has been evasive. The results above provide clear evidence of the effects of liquidity. This evidence can be used to assess the extent to which estimates of demand responses reflect the relaxation of liquidity constraints versus a (potentially inefficient) substitution effect.

Recent work by Ericson and Sydnor (2018) suggests that health insurance plan choices typically viewed as "mistakes" (Abaluck and Gruber, 2011) can actually be rationalized when considering liquidity constraints. Other work has also shown that consumers with non-linear health insurance contracts often respond to (seemingly irrelevant) current "spot" prices rather than the (seemingly relevant) expected end-of-year price (Aron-Dine et al., 2015; Einav et al., 2015). These spot-price responses can potentially reflect consumer mistakes or myopia, but they can also be rationalized by liquidity constraints (Ericson and Sydnor, 2018). Our results suggest that liquidity constraints may indeed play a major role in explaining both of these seemingly irrational consumer behaviors, given that liquidity constraints seem to have important effects on the consumption of prescription drugs, even for seniors with fairly generous insurance.

In summary, liquidity matters for healthcare consumption among the elderly. All in all, our results suggest that changes to the way health insurance is administered may improve the welfare of liquidity-constrained consumers. Future research ought to explore the direct effects of providing consumers with financing options or alternative ways of collecting copayments. Other forms of healthcare consumption, such as inpatient and outpatient hospital care, often involve ex-post out-of-pocket payments with limited consequences for delaying payments in the short-run. A similar structure for collecting drug copayments (ex-post instead of up-front) could relax the liquidity constraints we observe in Medicare Part D and allow seniors to cease delaying needed healthcare consumption until they once again find themselves with liquid assets.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210.

Agarwal, Sumit and Wenlan Qian, "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, 2014, *104* (12), 4205–4230.

_ , Chunlin Liu, and Nicholas S Souleles, "The Reaction of Consumer Spending and Debt to Tax Rebates Evidence from Consumer Credit Data," *Journal of Political Economy*, 2007, *115* (6), 986–1019.

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?," *Review of Economics and Statistics*, 2015, *97* (4), 725–741.

_ , _ , and _ , "The RAND Health Insurance Experiment, Three Decades Later," *Journal of Economic Perspectives*, 2013, *27* (1), 197–222.

Arrow, Kenneth J., "Uncertainty and The Welfare Economics of Medical Care," *American Economic Review*, 1963, *53* (5), 941–973.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *Quarterly Journal of Economics*, 2015, *130* (4), 1623–1667.

Baker, Scott R. and Constantine Yannelis, "Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown," *Review of Economic Dynamics*, 2017, *23* (1), 99–124.

Baugh, Brian and Jialan Wang, "When Is It Hard to Make Ends Meet?," *Mimeo*, 2018.

_ , Itzhak Ben-David, and Hoonsuk Park, "Disentangling Financial Constraints, Precautionary Savings, and Myopia: Household Behavior Surrounding Federal Tax Returns," Working Paper 19783, National Bureau of Economic Research 2014.

Beatty, Timothy K.M., Marianne P. Bitler, Xinzhe Huang Cheng, and Cynthia van der Werf, "SNAP and Paycheck Cycles," *Southern Economic Journal*, 2019, *86* (1), 18–48.

Berniell, Inés, "Pay Cycles: Individual and Aggregate Effects of Paycheck Frequency," *Mimeo*, 2018.

Bertrand, Marianne and Adair Morse, "What Do High–Interest Borrowers Do with Their Tax Rebate?," *American Economic Review*, 2009, *99* (2), 418–423.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "What Does a Deductible Do? The Impact of Cost–Sharing on Health Care Prices, Quantities, and Spending Dynamics," *Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.

Card, David, Raj Chetty, and Andrea Weber, "Cash–on–Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market," *Quarterly Journal of Economics*, 2007, *122* (4), 1511–1560.

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang, "Poverty and Economic Decision–Making: Evidence from Changes in Financial Resources at Payday," *American Economic Review*, 2016, *106* (2), 260–284.

Centeno, Mário and Álvaro A. Novo, "Reemployment Wages and UI Liquidity Effect: A Regression Discontinuity Approach," *Portuguese Economic Journal*, 2009, *8* (1), 45–52.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight, "Patient cost-sharing and hospitalization offsets in the elderly," *American Economic Review*, 2010, *100* (1), 193–213.

Cheng, Xinzhe Huang and Timothy K.M. Beatty, "Food Price Variation over the SNAP Benefit Cycle," *Mimeo*, 2016.

Chetty, Raj, "Moral Hazard vs. Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 2008, *116* (2), 173–234.

_ **and Amy Finkelstein**, "Social Insurance: Connecting Theory to Data," in Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *Handbook of Public Economics, Volume 5*, Elsevier, 2013, pp. 111–193.

Dobkin, Carlos and Steven L. Puller, "The Effects of Government Transfers on Monthly Cycles in Drug Abuse, Hospitalization and Mortality," *Journal of Public Economics*, 2007, *91* (11), 2137–2157.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf, "The Response of Drug Expenditure to Non–Linear Contract Design: Evidence from Medicare Part D," *Quarterly Journal of Economics*, 2015, *130* (2), 841–899.

_ **and** _ , "Moral Hazard in Health Insurance: What We Know and How We Know It," *Journal of the European Economic Association*, 2018, *16* (4), 957–982.

Ericson, Keith Marzilli and Justin R. Sydnor, "Liquidity Constraints and the Value of Insurance," Working Paper 24993, National Bureau of Economic Research 2018.

Evans, William N. and Timothy J. Moore, "The Short–Term Mortality Consequences of Income Receipt," *Journal of Public Economics*, 2011, *95* (11–12), 1410–1424.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and the Oregon Health Study Group, "The Oregon Health Insurance Experiment: Evidence from the First Year," *Quarterly Journal of Economics*, 2012, *127* (3), 1057–1106.

Foley, C. Fritz, "Welfare Payments and Crime," *Review of Economics and Statistics*, 2011, *93* (1), 97–112.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis, "How Individuals Respond to a Liquidity Shock: Evidence from the 2013 Government Shutdown," *Journal of Public Economics*, 2019, *Forthcoming.*

Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel, "Issuance and Incidence: SNAP Benefit Cycles and Grocery Prices," *Mimeo*, 2019.

Gross, David B. and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics*, 2002, *117* (1), 149–185.

Gross, Tal and Jeremy Tobacman, "Dangerous Liquidity and the Demand for Health Care Evidence from the 2008 Stimulus Payments," *Journal of Human Resources*, 2014, *49* (2), 424–445.

_ , Matthew J. Notowidigdo, and Jialan Wang, "The Marginal Propensity to Consume Over the Business Cycle," *American Economic Journal: Macroeconomics*, 2019, *Forthcoming.*

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 1997, *87* (1), 192–205.

Hansen, Gary D. and Ayse Imrohoroglu, "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard," *Journal of Political Economy*, 1992, *100* (1), 118–142.

Hastings, Justine and Ebonya Washington, "The First of the Month Effect: Consumer Behavior and Store Responses," *American Economic Journal: Economic Policy*, 2010, *2* (2), 142–162.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Kaiser Family Foundation, "Employer Health Benefits," http://files.kff.org/attachment/Report-Employer-Health-Benefits-Annual-Survey-2018 2018.

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

LaLumia, Sara, "The EITC, Tax Refunds, and Unemployment Spells," *American Economic Journal: Economic Policy*, 2013, *5* (2), 190–223.

Leary, Jesse B. and Jialan Wang, "Liquidity Constraints and Budgeting Mistakes: Evidence from Social Security Recipients," *Mimeo*, 2016.

Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, and Arleen Leibowitz, "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *American Economic Review*, 1987, *77* (3), 251–277.

Nyman, John A., "The Value of Health Insurance: The Access Motive," *Journal of Health Economics*, 1999, *18* (2), 141–152.

O'Carroll, Patrick, "Hearing on the Direct Deposit of Social Security Benefits," 2012.

Olafsson, Arna and Michaela Pagel, "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *The Review of Financial Studies*, 2018, *31* (11), 4398–4446.

Parker, Jonathan A., "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 1999, *89* (4), 959–973.

_ , Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–2553.

Pauly, Mark V., "The Economics of Moral Hazard: Comment," *American Economic Review*, 1968, *58* (3), 531–537.

Shapiro, Jesse M., "Is there a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 2005, *89* (2), 303–325.

Shigeoka, Hitoshi, "The Effect of Patient Cost Sharing on Utilization, Health, and Risk Protection," *American Economic Review*, 2014, *104* (7), 2152–2184.

Souleles, Nicholas S., "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 1999, *89* (4), 947–958.

Stephens, Melvin Jr., "'3rd of tha Month:' Do Social Security Recipients Smooth Consumption Between Checks?," *American Economic Review*, 2003, *93* (1), 406–422.

_ , "Paycheque Receipt and the Timing of Consumption," *Economic Journal*, 2006, *116* (513), 680–701.

Thaler, Richard H., "Mental Accounting and Consumer Choice," *Marketing Science*, 1985, *4* (3), 177–266.

_ , "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 1999, *12* (3), 183–206.

Vellekoop, Nathanael, "Explaining Intra–Monthly Consumption Patterns: The Timing of Income or the Timing of Consumption Commitments?," Working Paper 237, SAFE 2018.

Wright, Bill J., Matthew J. Carlson, Tina Edlund, Jennifer DeVoe, Charles Gallia, and Jeanene Smith, "The Impact of Increased Cost Sharing on Medicaid Enrollees," *Health Affairs*, 2005, *24* (4), 1106–1116.

Zeckhauser, Richard, "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives," *Journal of Economic Theory*, 1970, *2* (1), 10–26.

Zhang, C. Yiwei, "Consumption Responses to Pay Frequency: Evidence from 'Extra' Paychecks," *Advances in Consumer Research*, 2017, *45*, 170–174.

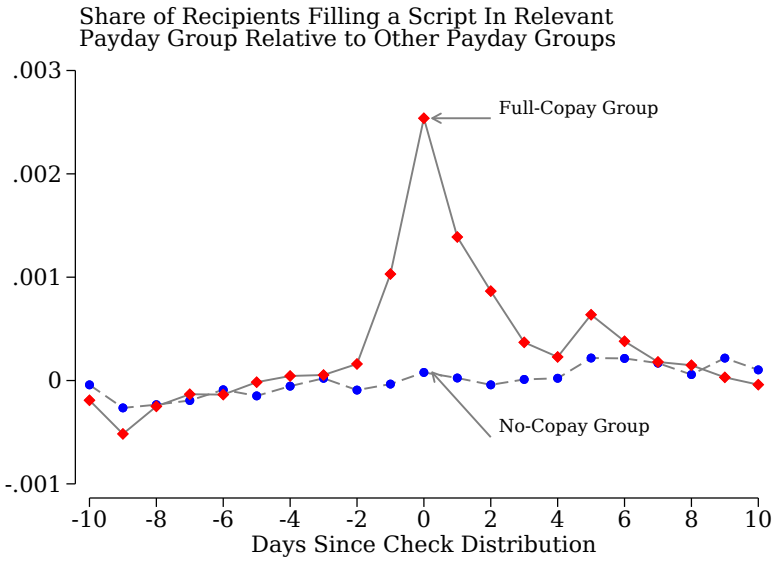
A Appendix Tables

[Under Construction.]

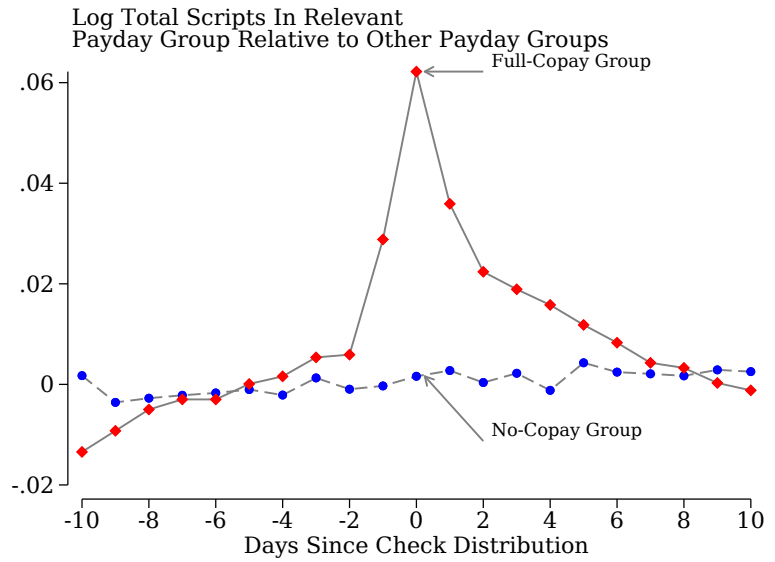
B Appendix Figures

Appendix Figure A1. Re-Centered Time Series

(a) Share of Beneficiaries Who Fill a Script By Day



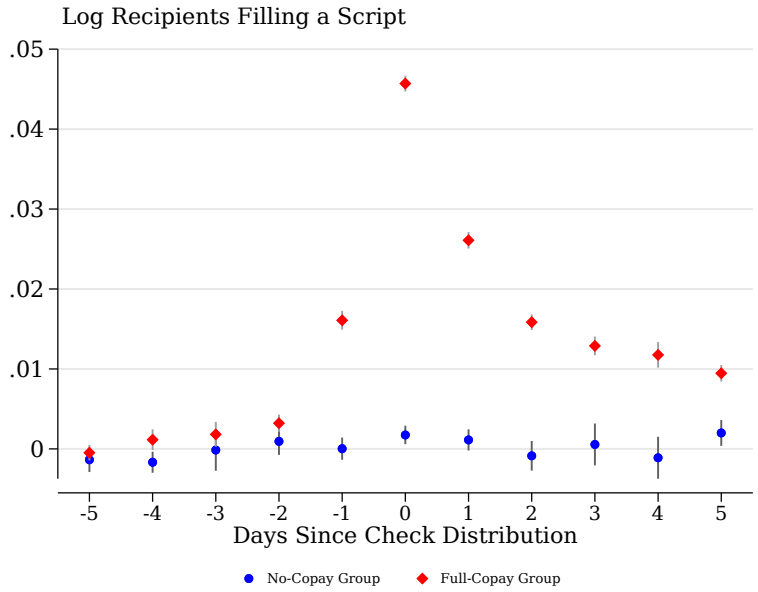
(b) Logarithm of Total Scripts



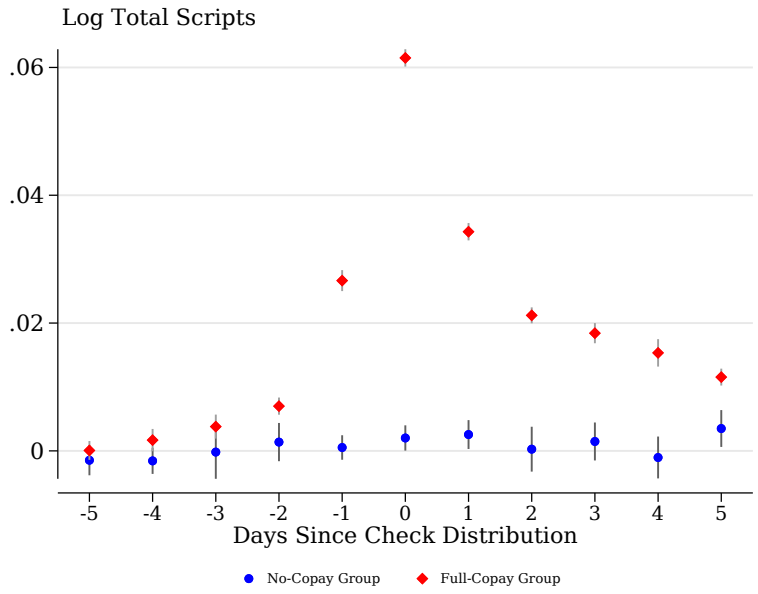
Figures plot measures of prescription drug consumption on each event-time day relative to check receipt for the no-copay group and the full-copay group. Panel (a) plots the share of beneficiaries who fill a script. Panel (b) plots log total scripts filled. For each copay group and relative day, we calculate the difference in fills between the beneficiaries who are scheduled to receive benefits versus the beneficiaries who are scheduled to receive their benefits with a two-week difference. For more details, see Section 3.1.

Appendix Figure A2. Event Study

(a) Log Beneficiaries Who Fill a Script By Day



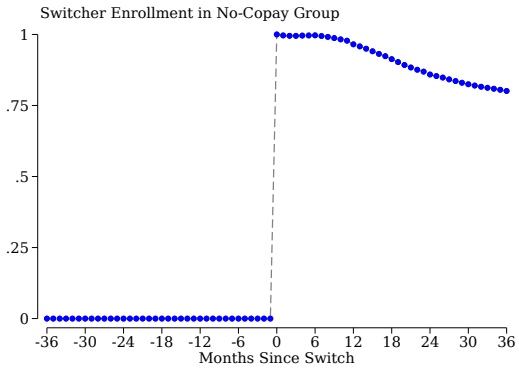
(b) Logarithm of Total Scripts



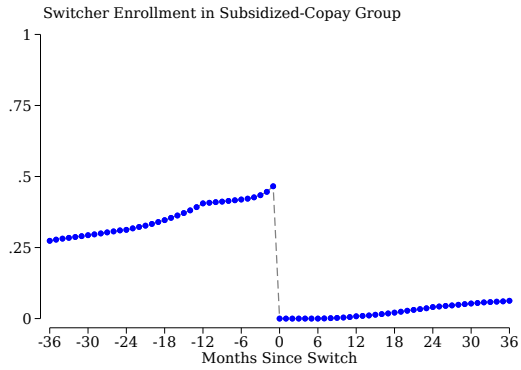
Figures plot event study estimates of changes in measures of prescription drug consumption on each event-time day relative to check receipt for the no-copay group and the full-copay group. Panel (a) plots estimates for the share of beneficiaries who fill a script. Panel (b) plots estimates for log total scripts filled. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details, see Section 3.2.

Appendix Figure A3. Switchers: Re-Centered Time Series of Enrollment

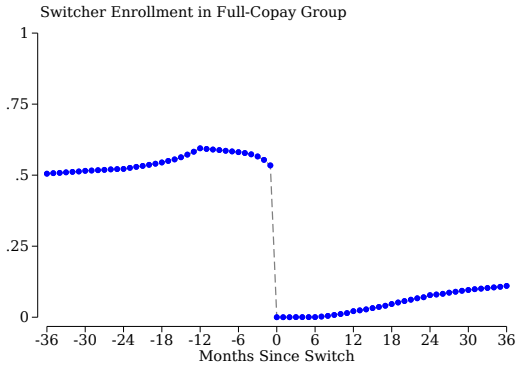
(a) No-Copay Group



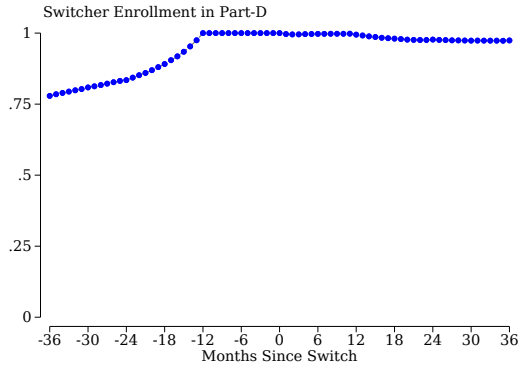
(b) Subsidized-Copay Group



(c) Full-Copay Group



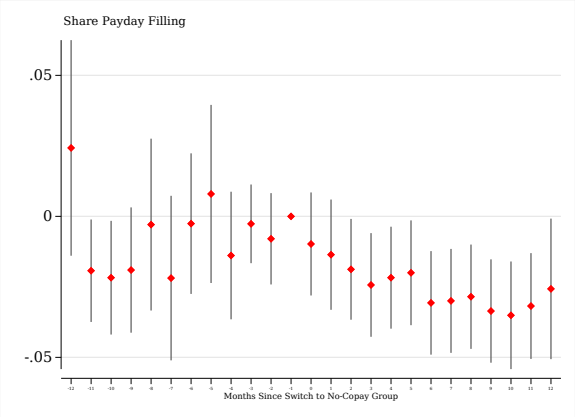
(d) Part-D



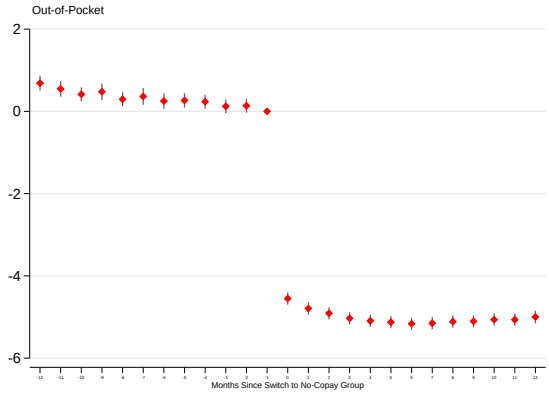
Figures plot the enrollment patterns of beneficiaries who switch to the no-copay group over time relative to the time of switching. Panel (a) shows enrollment in the no-copay group. Panel (b) shows enrollment in the subsidized-copay group. Panel (c) shows enrollment in the full-copay group. Panel (d) shows enrollment in Medicare Part D. For more details, see Section 4.1.

Appendix Figure A4. Switchers: Event Study

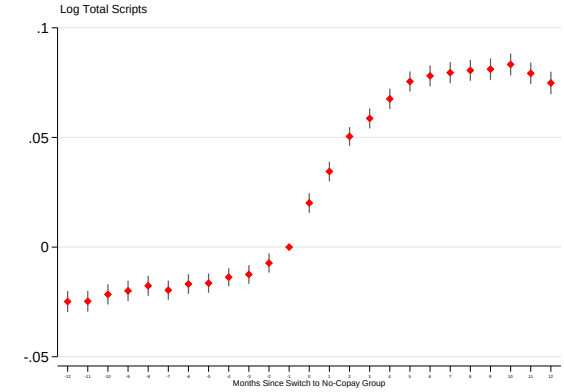
(a) Share Payday Filling



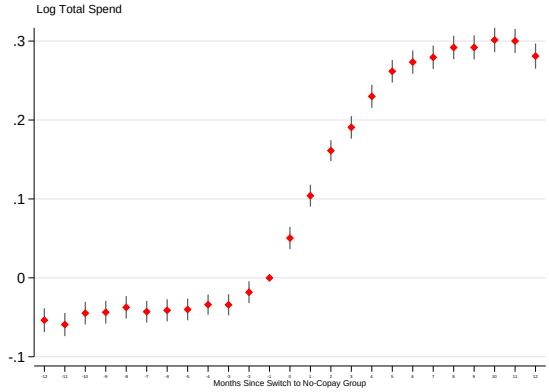
(b) Portion Out-of-Pocket



(c) Log Total Scripts



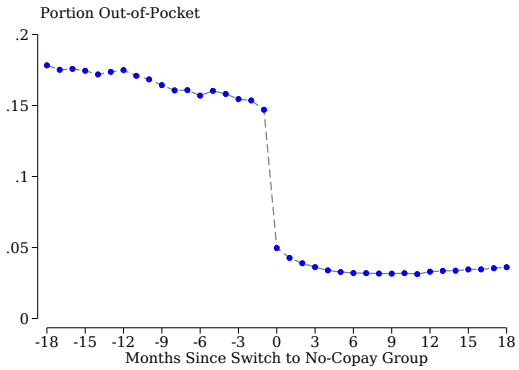
(d) Log Total Spending



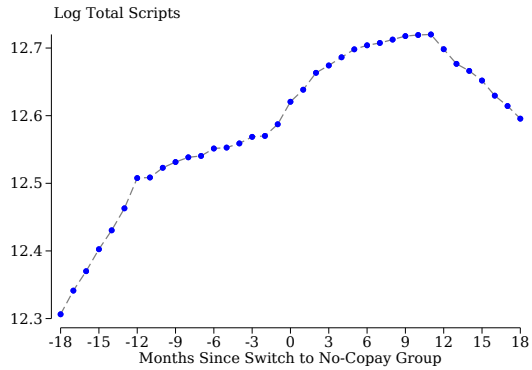
Figures show event study estimates of the impact of switching to the no-copay group on a variety of outcomes. Panel(a) shows the share payday filling. Panel (b) shows the portion of spending that is out-of-pocket. Panel (c) shows log total scripts. Panel (d) shows log total spending. For more details, see Section 4.1.

Appendix Figure A5. Switchers: Re-Centered Time Series

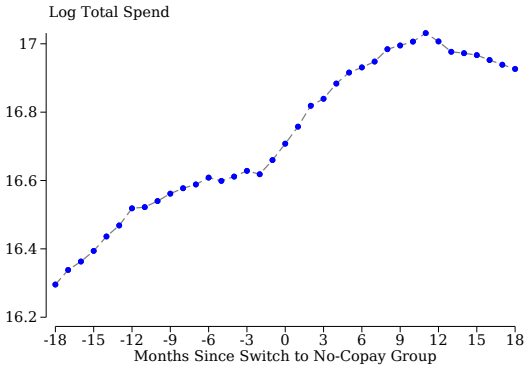
(a) Portion Out-of-Pocket



(b) Log Total Scripts



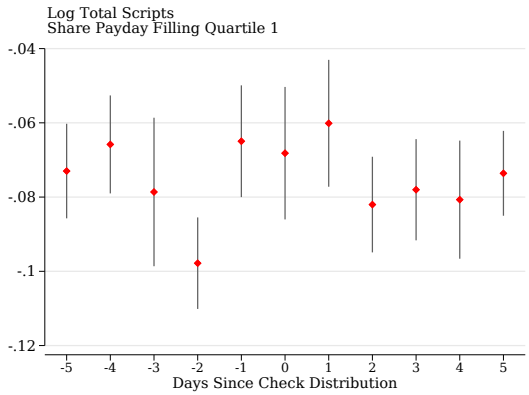
(c) Log Total Spending



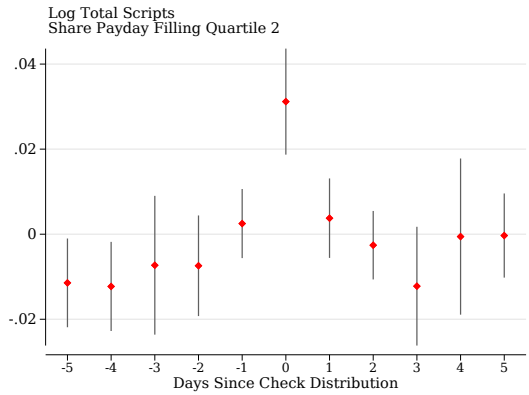
Figures show changes in measures of prescription drug expenditures to switching to the no-copay group. Panel (a) shows the portion of spending that is out-of-pocket. Panel (b) shows log total scripts. Panel (c) shows log total spending. For more details, see Section 4.1.

Appendix Figure A6. Logarithm of Total Scripts

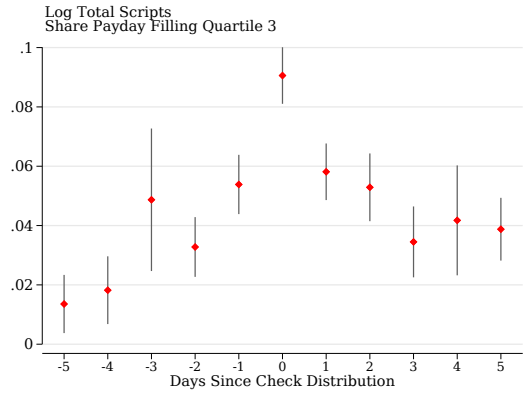
(a) Share Payday Filling Quartile 1



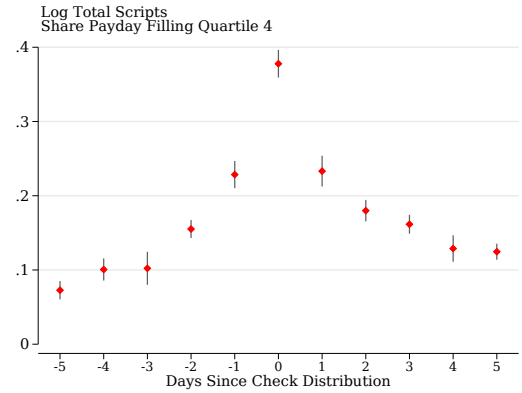
(b) Share Payday Filling Quartile 2



(c) Share Payday Filling Quartile 3



(d) Share Payday Filling Quartile 4



Figures plot event study estimates of changes in log total scripts filled on each event-time day relative to check receipt for the subsidized-copay group by quartile of the share of Wednesday fills that occur on payday. We estimate this share as a leave-one-out measure, calculating it over all years except the year for which we estimate the event study. These coefficients are from estimating Equation (1). Standard errors are clustered at the level of the birthday group. For more details on the share of Wednesday fills that occur on payday measure, see Section 4.1. For more details on the event study estimation, see Section 3.2.