THE
QUARTERLY JOURNAL
OF ECONOMICS

Vol. 139 2024 Issue 4

THE HEALTH COSTS OF COST SHARING *

AMITABH CHANDRA
EVAN FLACK
ZIAD OBERMEYER

What happens when patients suddenly stop their medications? We study the health consequences of drug interruptions caused by large, abrupt, and arbitrary changes in price. Medicare's prescription drug benefit as-if-randomly assigns 65-year-olds a drug budget as a function of their birth month, beyond which out-of-pocket costs suddenly increase. Those facing smaller budgets consume fewer drugs and die more: mortality increases 0.0164 percentage points per month (13.9%) for each $100 per month budget decrease (24.4%). This estimate is robust to a range of falsification checks and lies in the 97.8th percentile of 544 placebo estimates from similar populations that lack the same idiosyncratic budget policy. Several facts help make sense of this large effect. First, patients stop taking drugs that are both high value and suspected to cause life-threatening withdrawal syndromes when stopped. Second, using machine learning, we identify patients at the highest risk of drug-preventable adverse events. Contrary to the predictions of standard economic models, high-risk patients (e.g., those most likely to have a heart attack) cut back more than low-risk patients on exactly those drugs that would benefit them the most (e.g., statins). Finally, patients appear unaware of these risks. In a survey of 65-year-olds, only one-third believe that stopping their drugs for up to

* We acknowledge funding from the NSF Graduate Research Fellowship (EF), National Institute on Aging grant P01AG005842 (ZO), and the Chan Zuckerberg Biohub–San Francisco (where ZO is an Investigator). We thank, without implicating, Ned Augenblick, Stefano DellaVigna, Amy Finkelstein, Ben Handel, Peter Hull, Jon Kolstad, Jon Gruber, Robin McKnight, Sendhil Mullainathan, Mohan Ramanujan, Julian Reif, Jon Skinner, Josh Schwartzstein, Amar Venugopal, Anna Zink, five thoughtful referees, and numerous seminar participants for their comments.

The Quarterly Journal of Economics (2024), 2037–2082. https://doi.org/10.1093/qje/qjae015.
Advance Access publication on May 3, 2024.

a month could have any serious consequences. We conclude that far from curbing waste, cost sharing is itself highly inefficient, resulting in missed opportunities to buy health at very low cost ($11,321 per life-year). *JEL codes:* I12, I13.

I. Introduction

Randomized trials quantify the benefits of starting a drug, but we know far less about the costs of stopping a drug. A researcher interested in this topic would struggle to obtain ethical approval: experimentally withdrawing prescribed medications is infeasible in most settings. But the nonlinear contract structure of many health insurance plans causes millions of patients to stop prescriptions abruptly every year, due to sharp variation in out-of-pocket drug prices over time (Einav, Finkelstein, and Schrimpf 2015; Einav and Finkelstein 2018; Einav, Finkelstein, and Polyakova 2018).

In general, these price-driven drug interruptions have not been a cause for concern, at least to economists, because their welfare effects are theoretically ambiguous. Yes, patients stop taking apparently high-value medicines, like statins and beta-blockers (Choudhry et al. 2011; Brot-Goldberg et al. 2017; Einav, Finkelstein, and Polyakova 2018), but there is no clear-cut evidence that this harms their health.[1] Patients may be deciding, rationally, that financial savings outweigh health gains, particularly if they have private information about the benefits (or side effects) of treatment. Indeed, patients themselves seem unbothered. In a survey of Medicare-age patients taking medication (Figure I), we find that only one-third (33.5%) believe that missing their drugs for up to a month could have serious consequences (e.g., hospitalization, death). A majority (53%) predict no negative health consequences, even simply feeling worse, from missing their drugs for up to a week.

The medical literature, by contrast, gives more cause for concern. Beyond the obvious—stopping a drug means forgoing any benefits shown in clinical trials—interrupting some drugs can

1. Previous studies show only effects on proxies for health, like hospitalizations or spending. For example, Chandra, Gruber, and McKnight (2010) find that price increases in costs for drugs led to increased hospitalizations, but they did not study mortality.

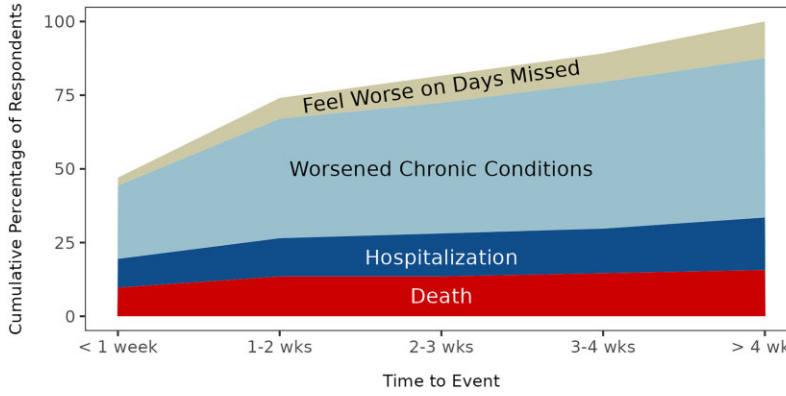

FIGURE I

Patient Beliefs about the Health Effects of Drug Interruptions

Results of a survey of Medicare-age patients taking at least one prescribed medication. The *y*-axis shows the cumulative percent who believed a given type of health problem could result from interrupting consumption of their medications. Respondents were able to choose multiple problems, so each individual is assigned the most severe problem, ordered from feeling worse to death. The *x*-axis shows the minimum number of weeks respondents thought it would take for such a health problem to occur. More details on the survey are in Section IV.

cause withdrawal or "rebound" effects.[2] This phenomenon is difficult to study experimentally, but there are some rare, striking signals from the medical literature that even short interruptions can have life-and-death consequences. Consider a classic study that randomly alternated a drug for high blood pressure (the beta-blocker propranolol) with placebo, to identify an optimal dose (Miller et al. 1975). Of 20 patients participating in what should have been a fairly mundane dose-finding trial, 10 had adverse cardiac events, 6 of which were heart attack or death; all adverse events occurred in the short placebo periods—4 of 44 total weeks—when the drug was abruptly withdrawn. In another study, researchers noted that patients taking statins sometimes failed to get them for idiosyncratic reasons, for example, in the

2. Rebound effects result from a set of equilibrating forces in the body that build up over time to compensate for a drug's effect (e.g., opioids gradually desensitize opioid receptors, requiring escalating doses to get the same effect; Dumas and Pollack 2008). When the drug is suddenly stopped, these endogenous mechanisms are no longer balanced by the drug, producing an opposite effect that will be problematic if the drug was doing something helpful.

setting of hospitalization for heart attack (Heeschen et al. 2002). Those who failed to receive their statin had threefold higher risk of death or repeat heart attack than similar patients who got their statin, and a 69% higher risk than patients who never took statins. Rebound effects have been suggested for a large number of drugs in wide use by older patients, and guidelines suggest tapering under medical supervision rather than simply stopping them (Steinman and Reeve 2023; Bain et al. 2008). But the empirical evidence backing these guidelines is weak, relying mostly on animal models or highly confounded real-world studies that simply compare patients who do versus do not interrupt their medications. Making matters worse, physicians are unlikely to notice even large effects: detecting a 30% increase in mortality from a baseline of 1% would require perfect awareness of drug interruptions in a sample size of 40,000 (versus a typical panel: 2,000 patients; Raffoul et al. 2016). So in practice, even if doctors advise against interruptions, their recommendations are lost on patients—as evidenced by our survey results.

We study the mortality effects of drug interruptions at scale, exploiting an abrupt and as-good-as-random price shock affecting millions of Medicare Part D beneficiaries. We build on an identification strategy pioneered by Aron-Dine et al. (2015) and Kaplan and Zhang (2017): a nonlinear contract structure causes drug prices to vary by birth month for enrollees in their first year of coverage. Historically, every January, beneficiaries start by paying only 25% of drug costs out of pocket. If they exceed an annual budget cap of ~$2,500, however, they enter the "donut hole," where out-of-pocket costs jump to 100%. Whether a beneficiary enters the donut hole is of course not random: it depends on prior consumption. Critically, however, the cap is not prorated, giving all enrollees the same "pre-donut budget" whether they enroll early or late in the year. By the time a later enrollee spends her first dollar, an earlier enrollee has been spending down her budget for months, approaching the donut hole. Because Part D eligibility begins in the month someone turns 65, birth month generates exogenous budget variation in the first year of enrollment.[3]

3. Reassuringly, baseline characteristics are similar for early versus late enrollees, and are balanced on demographics, drug consumption, and mortality in the first 90 days of enrollment.

An important nuance is that the donut hole affects beneficiaries differently, depending on their drug spending. For example, most people end the year with far less than $2,500 of total spending, meaning even early enrollees do not risk entering the donut hole. We focus on a set of middle spenders, whose initial drug spending puts them on track to enter the donut hole if they enroll early in the year but not if they enroll late.[4] We focus our analysis on the month of December, when our power to detect an effect is highest: sample size increases over time as 45,000 more beneficiaries enroll every month, and drug consumption differences across enrollment months grow as more early enrollees enter the donut hole.

The complex interaction of budget caps, initial drug spending, and enrollment month produces sharp variation in mortality. Enrolling one month earlier increases December mortality by 0.0112 percentage points, or 9.5%. We translate this into policy-relevant terms by noting that enrollment month implies a monthly drug budget before full cost sharing sets in: each $100/month drug budget decrease (a 24.4% change versus the average budget) increases mortality by 13.9%.[5] We verify that changes in consumption mirror changes in assigned drug budgets and mortality, with earlier enrollees consuming significantly fewer drugs.[6] The large effects in middle spenders contrast with the absence of any effects in low spenders, just as we expect: low spenders remain far from the donut hole, even if they enroll

4. We verify that initial spending, like other pretreatment characteristics, is uncorrelated with enrollment month. This means specifically that forward-looking behavior has not yet induced differential selection into initial spending bins, allowing us to identify a similar group of middle spenders across enrollment months. We consider high spenders, who not only enter the donut hole but also approach or enter the catastrophic coverage, in detail in Section II.C.

5. Monthly budget is a mechanical function of enrollment month: for example, a $2,500 budget gives a February enrollee $227/month and a September enrollee $625/month.

6. We cannot easily link mortality to drug consumption using 2SLS in our setting, as the intertemporal dependence of drug consumption and mortality would introduce bias. Intuitively, the enrollment-month effect on mortality is mediated via a feedback loop of prices and quantities over time. Using any one of these quantities as the endogenous variable would violate the exclusion restriction. In particular, Online Appendix B shows that the combination of (i) prior drug consumption effects on current mortality and (ii) intertemporal substitution of drugs across periods produces a large upward bias in the estimated effect of consumption on mortality in any one period.

early, and show no effect of enrollment month on consumption or mortality. At the same time, the highest-spending 2%–3% show large, significant, and opposite-sign effects of enrollment month versus middle spenders. Again, this is just as we expect: in this group, earlier enrollees spend through the donut hole and enter the "catastrophic coverage," where cost sharing drops from 100% to near zero, resulting in more drug consumption and lower mortality; later enrollees remain stuck in the donut hole. To summarize, enrollment month has effects on both drug consumption and mortality that vary widely across subpopulations, in a way that exactly matches the interaction of Part D budget caps and initial spending. It is hard to imagine a confounder with such an idiosyncratic structure.

To build confidence in our estimates, we conduct a broad set of falsification tests and checks, beginning with two obvious potential confounders. First, while earlier enrollees are by construction slightly older, mortality differences are 10 times larger than those implied by age differences (based on Social Security data). Second, well-known health differences across birth seasons have a different temporal pattern and are far too small to explain our results (Doblhammer and Vaupel 2001). Next, we replicate our analysis in a wide range of closely related settings, to ensure that mortality effects are only present when enrollment month affects drug budgets. For example, replicating our analysis before versus after reforms to Medicare that began to close the donut hole in 2011, we find that the effect size mirrors the extent of cost sharing. We also track the evolution of enrollment-month effects on drug consumption and mortality in our cohort of middle spenders, before and after December of the first year. We find a significant difference in mortality only in December, exactly when differences in consumption across enrollment months peak. The effect fades, then disappears after prices reset for all enrollees in January. Finally, we formally generate a large number of placebo estimates: enrollment-month effects on mortality, in Medicare populations for whom enrollment month does not affect drug budgets. This includes, for example, 66-year-olds, who are no longer affected by the enrollment-month quirk we exploit at age 65; 65-year-old dual-eligibles who face no cost sharing; and 64-year-old disabled beneficiaries whose enrollment timing is not driven by birth month. Replicating our analysis in each of these samples, across a range of calendar months, our main estimate lies in the 97.8th percentile of mortality effects, larger in absolute value than 532 of 544 placebo estimates.

We make sense of this large mortality effect in two ways. First, like Einav, Finkelstein, and Polyakova (2018) and Brot-Goldberg et al. (2017), we find that patients interrupt drugs that are high value. Many are also suspected to produce dangerous rebound syndromes when interrupted, with potentially life-threatening consequences: statins, antihypertensives, glucose-lowering agents, inhalers, and steroids for pulmonary disease. Second, among those taking a specific drug, we document a positive correlation between treatment benefit and likelihood of price-driven interruption. For example, beneficiaries at the highest risk of heart attack and stroke cut back four times as much on cardiovascular drugs (e.g., statins, antihypertensives) versus lower-risk patients (2.65 versus 0.766 drug-days per $100 change in monthly pre-donut coverage, a 4.0% versus 1.8% reduction).[7] Similar patterns exist for diabetes and respiratory drugs. These differences are unlikely to be explained by income alone, as we see similar cutbacks in both high- and low-income ZIP codes. One potential driver of this effect is that higher-risk patients are more likely to interrupt all of their medications when prices increase, causing large absolute reductions in those on more drugs at baseline.

Our findings may be surprising from the point of view of standard economic models of behavior that emphasize moral hazard or private information: patients should not interrupt drugs with large benefits. Yet they fit with a growing literature in economics linking insurance coverage to lower mortality. Most closely related, Abaluck et al. (2021) find that switching into a plan with donut-hole coverage reduces mortality by 9.8%.[8] More broadly, Miller, Johnson, and Wherry (2021) show Medicaid expansion reduced mortality by 9.4%, and Goldin, Lurie, and McCubbin (2021) find that Affordable Care Act tax incentives reduced mortality by 6.3%. Together, these studies show that health insurance can have large health effects by affecting how patients use—and especially underuse—high-value care (Baicker, Mullainathan, and Schwartzstein 2015).

7. There is strong evidence from the medical literature that such risk is a good proxy for treatment benefit, particularly for cardiovascular drugs, an assumption we discuss in detail in Section III.D. Again, we caution against comparing the magnitudes of changes in drug-days and mortality, given the biases detailed in Online Appendix B.

8. This estimate is from a 2SLS with donut-hole coverage as the endogenous variable; of course, many other aspects of plans also vary so these comparisons are necessarily approximate. Further discussion of how our estimates relate to this literature is in Online Appendix Table E.2.

Cost sharing has been a cornerstone of health insurance design for decades, driven by worries about wasteful spending and moral hazard. Our results indicate that far from reining in low-value care, cost sharing is itself highly wasteful. Eliminating cost-driven drug interruptions would extend life at a cost of around $11,321 per life-year (versus commonly used thresholds for cost-effectiveness of $100,000–$200,000 per life-year; Neumann, Cohen, and Weinstein 2014).[9] An optimistic view of these results is that policy makers have a unique opportunity to purchase health at negligible cost by eliminating cost sharing for high-value, high-risk prescription drugs.

II. Empirical Strategy

Studying the health effects of drug interruptions is difficult. Experimentally stopping prescribed medications is impractical and unethical, and observational comparisons are highly confounded: patients who interrupt consumption are very different from those who do not. Perhaps the best evidence of this comes from randomized trials, in which every drug dose is carefully tracked: patients who take a lower fraction of assigned doses have worse outcomes—whether they are in the treatment or the placebo group (Osterberg and Blaschke 2005). This fact elegantly demonstrates how difficult it is to study this phenomenon: the myriad unmeasured links between drug-taking behavior and health confound comparisons of patients who do versus do not interrupt their medications, even in the setting of a randomized trial. The idiosyncratic structure of the Medicare drug benefit, which creates exogenous and abrupt changes in drug prices, provides a unique setting in which to study the effects of drug interruptions on mortality.

II.A. Medicare Policy Context

Since 2006, Medicare Part D has offered prescription drug coverage to seniors and disabled individuals in the United States. Individuals can enroll in either stand-alone prescription drug plans (PDPs) that are offered alongside traditional Medicare, or

9. Given that drug consumption can also offset inpatient spending (Chandra, Gruber, and McKnight 2010), this number is likely to be an overestimate of the cost. Unfortunately, because half of our sample is in Medicare Advantage and we do not observe in- or outpatient spending, we are unable to study this directly.

Medicare Advantage (MA) plans where drug coverage is bundled with inpatient/outpatient care. The benefit's nonlinear structure with respect to out-of-pocket costs is illustrated in Figure II, Panel A (Einav, Finkelstein, and Schrimpf 2015).

Using plan details from 2008 to describe the plan, the calendar year begins with a deductible phase in which the beneficiary pays the entire cost of all drugs until spending reaches $275. She then faces a 25% cost-sharing (coinsurance) rate that lasts until spending exceeds the budget cap of $2,510 (the initial coverage limit). After this, the beneficiary falls into the coverage gap, or donut hole, and again pays 100% of the cost of all drugs (based on list prices, which are significantly higher than net prices). Finally, after reaching $5,726 of total spending, she enters the catastrophic coverage (to continue the analogy, this phase represents the far side of the donut). Here, she pays only 5% of the drug cost, or a copay of $2.25 to $5.60 for each drug (depending on whether it is generic or preferred, respectively).[10] The cutoff points for each coverage arm change slightly from year to year, as shown in Online Appendix Table A.1, but the basic structure remained the same until 2011, when the donut hole began to close as a result of policy changes: cost-sharing rates for generic and branded drugs in the donut hole fell from 100% to 50% and 93%, respectively. The donut hole was officially eliminated by the Affordable Care Act in 2020, but many aspects of its structure persist in both Medicare (CMS 2024) and private insurance.[11]

As first noted by Aron-Dine et al. (2015) and Kaplan and Zhang (2017), the spending limits that define cost sharing are not prorated in the first year of enrollment. In other words, a person who enrolls in, say, February gets the same pre–donut hole budget, $2,510 of total spending, as a person who enrolls in Septem-

10. The overall nonlinear structure was largely the result of a political compromise balancing the desire to cover very sick beneficiaries (the catastrophic phase) with reducing the cost of the program overall (the donut hole); a review is found in Oliver, Lee, and Lipton (2004). Insurers may offer coverage that is actuarially equivalent, or enhanced compared with the standard benefit. One common deviation from the standard design is to replace the deductible phase with uniform cost sharing until the donut hole is reached. In addition, most plans do not use coinsurance, but use copays based on drug tiers for each coverage arm. In practice, copays equate to roughly the same level of cost sharing in each arm as the coinsurance rates specified by the standard benefit (Einav, Finkelstein, and Schrimpf 2015).

11. For a helpful overview, see KFF (2023).

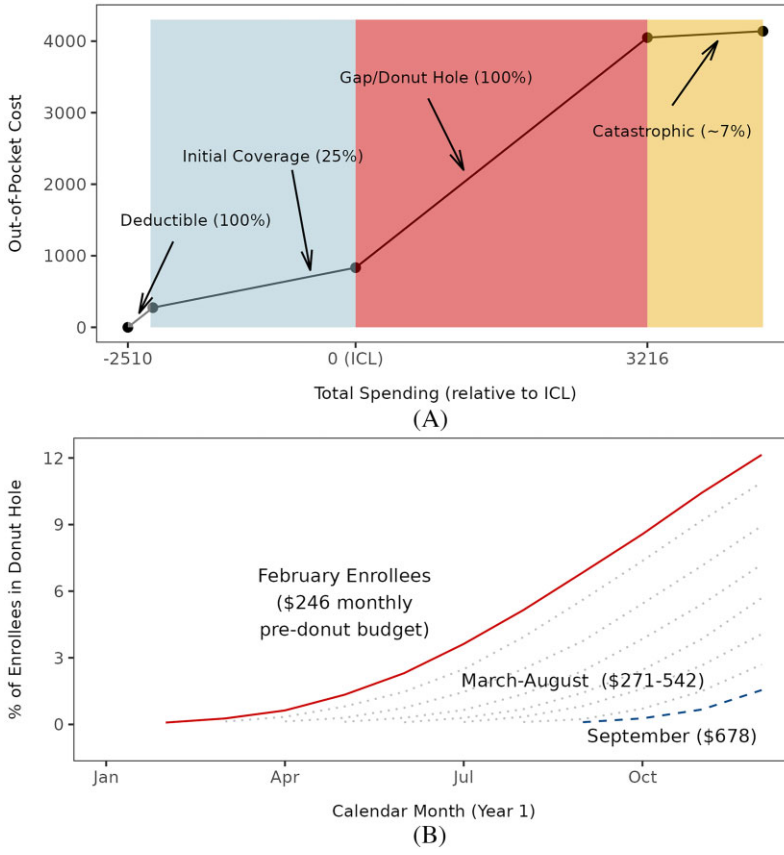

FIGURE II

Medicare Drug Benefit Design and Donut-Hole Entry

Panel A: Part D standard benefit design, adapted from Einav, Finkelstein, and Schrimpf (2015), using 2008 program details. The initial coverage limit (ICL) is the budget cap where beneficiaries transition from initial coverage to the donut hole. Panel B: Percentage of beneficiaries who enter the donut hole by the end of their first calendar year of enrollment, by enrollment month. February enrollees appear on top (solid red line; color version available online), March–August enrollees in the middle (dotted gray lines), and September enrollees at the bottom (dashed blue line). The monthly pre-donut budget, the amount each beneficiary can spend before entering the donut hole, is shown in parentheses beside the enrollment month.

ber. It is easy to see that a beneficiary's monthly budget will be far greater if they enroll in later calendar months: the same pre-donut budget of $2,510 is spread over fewer months. This results in earlier enrollees being more likely to exceed their budget cap, landing them in the higher prices of the donut hole, while later enrollees are more likely to remain in the generous initial coverage phase with lower prices.[12]

Individuals become eligible to enroll in Part D on the first day of their 65th birth month, meaning that enrollment month is primarily driven by birth month. If we believe that (i) birth month is as-good-as-random with respect to health, and (ii) different birth months select into enrollment months similarly,[13] then the effect of any resulting variation in drug budgets on mortality can be identified. Like Aron-Dine et al. (2015), we prefer estimates based on enrollment month as opposed to birth month and use these in our main specifications. Enrollment month more accurately measures the start of actual spending, as opposed to the start of eligibility, meaning it more accurately predicts drug prices.[14] In the empirical work, we test the underlying assumptions in detail: first, by checking that those from different birth or enrollment months are similar on a range of observable health and utilization measures. We also compare our estimates to those from the literature on birth-seasonal variation in health to see if these known effects can account for our results.

12. Very high-spending earlier enrollees are also more likely to enter the catastrophic coverage after the donut hole, for the same reason.

13. Unlike Medicare Part A, enrollment in Part D is not automatic. Beneficiaries can enroll during a seven-month-long initial enrollment period that runs from three months before to three months after their 65th birth month. Coverage starts on the first day of the month after the individual enrolls, but not before the first day of the enrollee's birth month. If an individual chooses to enroll later, she faces a penalty of higher premiums for the remainder of her tenure on Part D. We explore the selection of enrollment timing from birth months in Online Appendix Figure A.1. Of those enrolling within a year of turning 65, the majority enroll in their birth month (69%). Empirically, a small number of individuals enroll in the month before their birth month. This proportion is similar across birth months and of similar magnitude in Aron-Dine et al. (2015).

14. Using birth month instead of enrollment months for our main analysis yield similar, if less precise, results, just as we would expect: see Online Appendix Table C.2.

II.B. Model

We would ideally like to estimate the effect of drug consumption Q, a vector of consumption measures for individual drugs (specifically, number of drug-days consumed in some period), on mortality, denoted by indicator Y. In addition to Q, Y is determined by patient factors, which can be either observed (vector X) or unobserved (vector W). Unobserved factors affect both mortality and consumption, so OLS regression of Y on Q will be confounded, hence the need for an instrument Z that exogenously shifts drug consumption: enrollment month. For ease of exposition, we begin with a simplified scenario where Z takes two values, indicating whether patients are assigned to enroll early in their first year, receiving the usual Medicare monthly drug budget ($Z = 0$), or late in the year, receiving a higher monthly drug budget ($Z = 1$).[15]

We begin by considering the effect of Z on drug prices via the donut hole.[16] Under Medicare policy, beneficiary i enters the donut hole if her cumulative spending from enrollment month $Z_i = z$ to period t, $S_{i,z:t} = \sum_{j=z}^{t} S_{ij}$ exceeds the donut-hole limit θ_D.[17] Let indicator $D_{it} = \mathbb{1}\left\{S_{i,z:t} \geqslant \theta_D\right\}$ measure whether the beneficiary is in the donut hole in t. The instrument mechanically affects the calculation of cumulative spending, making earlier enrollees more likely to enter the donut hole: in potential-outcomes notation, $D_{it}^0 \geqslant D_{it}^1$, where the superscript corresponds to the two potential enrollment months. This sets the vector of individual drug prices $P_{it} = P_{it}^0 + \delta D_{it}$, the baseline price P_0 plus the donut-hole premium δ. An earlier enrollee thus faces higher prices, $P_{it}^0 > P_{it}^1$, leading to lower drug consumption $Q_{it}^0 < Q_{it}^1$. We hypothesize that these differences in consumption drive mortality differences such that $Pr(Y_{it}^0 = 1) > Pr(Y_{it}^1 = 1)$.

Given our interest in the effect of budget-induced drug consumption changes on mortality, a natural estimation approach

15. This simplified binary case sets up our potential-outcomes framework; it will generalize to our actual empirical setting, where Z runs from 2 for February enrollees to 9 for September enrollees, with larger values always implying larger pre–donut hole budgets.

16. The donut hole is the policy feature that affects the largest fraction of beneficiaries. We consider the catastrophic-coverage policy, which only the highest-spending 2%–3% of beneficiaries enter, in Section II.C.

17. To simplify notation, we write θ_D as a constant, but in our analysis we account for the changing values of all policy thresholds across years.

would be 2SLS, instrumenting for drug consumption using enrollment month. Unfortunately, any simple choice of endogenous variable for the first stage will violate the exclusion restriction: the instrument affects mortality via a complex feedback loop connecting quantities to prices over multiple time periods, and any of these may affect mortality. A pernicious consequence of this problem is that estimates of the effect of consumption on same-period mortality will be biased upward, as we show more formally in Online Appendix B. The intuition for this is twofold. First, because prior and current drug consumption are correlated and both affect mortality, a regression of mortality on same-period consumption will be biased: the effect of prior consumption will be misattributed to the current period.[18] Second, because of the well-documented phenomenon of intertemporal substitution across time periods (Aron-Dine et al. 2015), past consumption is on average greater than current-period consumption, inflating the bias: small changes in current consumption will appear to have large effects on same-period mortality, partly reflecting the large changes in prior consumption.

As a result, we opt for a transparent, reduced-form estimate of the effect of enrollment month on mortality in period t, using the following estimating equation:

$$Y_{it} = \gamma_0 + \gamma_1 Z_i + X_i \gamma_2 + \gamma_{year} + \gamma_{plan} + \epsilon_{it}, \tag{1}$$

where γ_{year} is a set of calendar year fixed effects, γ_{plan} is a set of fixed effects for Part D plans, X_i is a vector of sex and race indicators, and the instrument Z_i is a scalar that takes integer values from 2 (February enrollment) to 9 (September enrollment).[19] Our primary interest is γ_1, the effect of enrolling one month later. The reduced-form estimate of enrollment month is not generalizable to other (non-Medicare) settings, where enrollment month is not a meaningful quantity. However, it does have a natural policy-relevant interpretation: enrollment month mechanically implies a monthly dollar budget—how much a beneficiary can spend before reaching the full cost sharing of the donut hole: $B_i = \frac{\theta_D}{(T - Z_i + 1)}$,

18. We cannot easily get around this by using cumulative consumption, because the measurement of consumption (and all other variables) in prior periods is correlated with the instrument—enrollment month: we will not observe consumption for later enrollees as we go back in time.

19. We exclude January, October, November, and December enrollees for reasons described in Section II.D.

where $T = 12$ (the total number of months). This has direct relevance to other health insurance settings that use similar budget structures.

Although we do not implement formal 2SLS to link mortality to drug consumption, we estimate the instrument's effect on drug consumption, for individual drugs and total consumption, in a single period t, via

$$Q_{it} = \alpha_0 + \alpha_1 Z_i + X_i \alpha_2 + \alpha_{year} + \alpha_{plan} + \eta_{it}. \quad (2)$$

This is useful to confirm that budget limits produce the expected changes in drug consumption and to study how beneficiaries choose to cut back. However, we emphasize that it is misleading to correlate the magnitude of changes in estimated consumption to changes in estimated mortality: any such effort suffers from the same bias as described above and in Online Appendix B, where small changes in consumption will appear to have outsized effects on same-period mortality.

Given the intertemporal dynamics, the choice of t is important. We favor a time period near the end of the first calendar year—specifically December, like Aron-Dine et al. (2015)—for two reasons. First, more and more beneficiaries enroll and reach their steady-state drug consumption, meaning we have larger samples as the year progresses. Second, differences in drug consumption across enrollment months accelerate over the course of the year, as shown in Figure II, Panel B, before prices reset for all enrollees in January. The combination of these two factors means our power to detect an effect increases dramatically and nonlinearly over the course of the first calendar year, peaking in December, as shown in Online Appendix Figure C.3. For example, our power to detect an effect of the same magnitude as December would be 72% lower in November and 76% lower in October.

II.C. Heterogeneity in the Effect of Enrollment Month on Prices

We have so far focused on the dominant effect of enrollment month Z_i: earlier enrollment increases the likelihood of entering the donut hole by year-end ($D_{i,12}$), resulting in higher prices and lower consumption. But the effect of Z_i is heterogeneous. For example, after the first calendar year of enrollment (where all beneficiaries look like our early enrollees, since cumulative spending is calculated starting in January), the majority of Part D beneficiaries have cumulative year-end spending well below the donut-

hole budget limit $E[S_{i,1:12}] \ll \theta_D \approx \$2500$.[20] For such low spenders $i \in \mathcal{L}$, the effect of enrollment month on donut-hole entry is likely negligible, because even the earliest enrollees will end the year far from the donut-hole limit. By contrast, a handful of the highest-spending beneficiaries $\mathcal{H}$ are *less* likely to end the year in the donut hole: they enter the donut hole, but then exit it when they exceed the second policy limit $E[S_{i,1:12}|i \in \mathcal{H}] \geqslant \theta_C \approx \$6000$, landing in the catastrophic-coverage phase where cost sharing returns to very low levels.

We would ideally like to focus attention on a group of middle spenders $\mathcal{M}$ whose spending is "just right": they spend enough to be affected by the donut-hole budget limit by year-end but not enough to approach the catastrophic-coverage limit. In this group, we expect large effects of enrollment month on donut-hole entry by year-end, and thus consumption and mortality, in a way that is relatively easy to interpret. An analogy can help see why: the middle spenders are like compliers, for whom donut-hole entry is monotonically lower for later enrollees versus earlier enrollees: $E[D^0_{i,12}|i \in \mathcal{M}] > E[D^1_{i,12}|i \in \mathcal{M}]$. The lowest spenders are never-takers, for whom $E[D^0_{i,12}|i \in \mathcal{L}] = E[D^1_{i,12}|i \in \mathcal{L}] = 0$. The highest spenders are defiers, $E[D^0_{i,12}|i \in \mathcal{H}] < E[D^1_{i,12}|i \in \mathcal{H}]$, who violate the monotonicity assumption: early enrollees are more likely to enter the donut hole (resulting in higher prices), but also more likely to enter the catastrophic coverage (resulting in lower prices)—that is, the effect of enrollment month Z_i on prices and consumption is ultimately reversed. Focusing on the middle spenders, who experience only the (prospect of) budget decreases, produces the most straightforward estimate of the effect of Z_i.

To assign beneficiaries to groups $\mathcal{G}_i \in \{\mathcal{L}, \mathcal{M}, \mathcal{H}\}$, we cannot simply use realized year-end spending—our identification strategy relies on the fact that cumulative year-end spending $S_{i,\,z:12}$ is endogenous to Z_i. Just after enrollment, though, the donut-hole limit may be sufficiently far away that spending is temporarily unaffected by Z_i: enrollees are all just starting to learn about a highly complex program. While this is plausible, it must be verified. If initial spending is affected by forward-looking behavior and thus correlated with enrollment month, using initial spending to assign groups would induce differential sample selec-

20. This is based on data from after the first enrollment year, that is, absent enrollment month effects.

tion across enrollment months: higher-spending earlier enrollees, who have already started to cut back, would be inappropriately grouped with lower-spending later enrollees, who have not. Fortunately, this is an empirical question: we can directly inspect beneficiaries' spending trajectories in their first months of Part D coverage, to determine whether initial spending varies by enrollment month. To do so, we first test for differences in spending in the first month of enrollment $E[S_{i,z}|Z_i = 1] - E[S_{i,z}|Z_i = 0]$ by regressing $S_{i,z}$ on scalar Z_i. If we find no significant effect (at $p < .05$), it implies that forward-looking behavior has not materialized by the end of the first month. We proceed to test for differences in month $(z + 1)$, both individually and cumulative since enrollment, and so on until we find a significant coefficient on Z_i. Online Appendix Table A.2 shows that such a difference emerges only in month $z + 3$. So we define "initial spending" as spending in the first three months of enrollment $S_{i,z:z+2}$. We calculate percentile bins of initial spending for each enrollment month, and use these to define a vector of indicators $\sigma_i = [\sigma_{i1}, \sigma_{i2},..., \sigma_{ij},...\sigma_{i100}]$, which are set to 1 if beneficiary i falls into the jth percentile of initial spending within her enrollment month. Using within-enrollment-month percentile is attractive because it requires only a weaker assumption, that the ranking of initial spending within enrollment month is correlated with donut-hole entry by the end of the year.

This lets us estimate the empirical analog of $E[D_{it}^1 - D_{it}^0|\mathcal{G}_i]$ using our instrument and spending bins: $E[D_{it}|Z_i = 1, \sigma_i] - E[D_{it}|Z_i = 0, \sigma_i]$.[21] Focusing on December $(t = 12)$, we regress

$$D_{i,12} = \phi_0 + (Z_i \times \sigma_i)\phi_1 + X_i\phi_2 + \phi_{year} + \phi_{plan} + \omega_{i,12} \quad (3)$$

and inspect vector ϕ_1, coefficients for the spending-percentile indicators σ_i interacted with scalar Z_i. This measures the effect of enrollment month on donut-hole entry at year-end for each spending percentile bin j. If $\phi_{1j} = 0$, it implies $E[D_{it}|Z_i = 1, \sigma_i] = E[D_{it}|Z_i = 0, \sigma_i]$, and all beneficiaries in spending percentile j are assigned to $\mathcal{L}$. Likewise, if $\phi_{1j} < 0$ then $i \in \mathcal{M}$, and if $\phi_{1j} > 0$ then $i \in \mathcal{H}$. With spending groups defined, we run our primary estimating equation (1) separately by group:

$$Y_{i,12} = \gamma_0 + (Z_i \times \mathcal{G}_i)\gamma_1 + X_i\gamma_2 + \gamma_{year} + \gamma_{plan} + \epsilon_{i,12}. \quad (4)$$

21. We take this empirical approach, rather than trying to forecast whether they will exceed policy budget limits (θ_D, θ_C) based on initial spending, because we know cutbacks start well in advance of these limits.

The vector γ_1 will measure the effect of enrollment month on mortality, separately by initial spending group $\mathcal{G}$. Our hypothesis is that these mortality effects will mirror the effects on donut-hole likelihood from equation (3), with $\gamma_{1\mathcal{L}} = 0$, $\gamma_{1\mathcal{M}} < 0$, and $\gamma_{1\mathcal{H}} > 0$.

II.D. Data

Our main sample for estimation consists of a 20% random sample of first-time Medicare Part D enrollees in their initial enrollment period (birth month and three subsequent months) from 2007 to 2012. First, we make sample restrictions common to the Part D literature—though we later use some of these excluded populations for falsification tests. We subset to those who become eligible for Medicare at age 65, under the Old Age and Survivors Insurance, and exclude those who enroll before age 65 for disability or end-stage renal disease. This leaves us with 1,131,922 observations. We then remove all individuals dually eligible for Medicaid or other low-income subsidies, as they face low prices that do not change as a function of drug spending, which leaves us with 925,170 observations. We also remove all individuals that enroll in a deductible plan, as their initial claims vary with enrollment month due to the future price effects (Einav, Finkelstein, and Schrimpf 2015), bringing the sample to 605,502. A series of other minor subsets brings our sample to 557,999 beneficiaries.[22]

We make three additional exclusions with respect to the timing of enrollment month and death. First, to calculate mortality rates in December, we exclude those who die before December 1. We carefully check mortality differences across enrollment months before December 1, which could indicate selection bias introduced by this exclusion, and find none: see Table II. We also drop those who enroll in October and later, because as Aron-Dine et al. (2015) note, these beneficiaries are still ramping up their drug consumption. As a result, their December utilization is spuriously low compared to beneficiaries enrolling earlier in the year, who have reached steady state in terms of consumption by December. Finally, we follow Aron-Dine et al. (2015) and exclude January enrollees from our sample for several reasons. Those born in October and later are legally allowed to enroll in Jan-

22. We exclude individuals in special needs plans, those with nonstandard ICL locations, and those not residing in the U.S. 50 states or Washington, DC. We include individuals in standalone PDPs and standard MA plans that are HMO, HMO POS, local PPO, private FFS, and regional PPO.

uary without penalty, because January is in their four-month initial enrollment period (IEP). Empirically, January enrollment appears to be an outlier in terms of volume of patients enrolling, and January enrollees are observably different from all other enrollment months.[23] With these restrictions, our final analytic sample consists of 358,706 individuals.

Our analyses of drug consumption use prescription fill date, total cost, out-of-pocket cost, and National Drug Code (NDC) identifiers. To classify drugs into clinically meaningful categories, we use RxNorm and RxClass APIs to link NDCs to Anatomical Therapeutic Chemical (ATC) codes, a hierarchical system that allows us to map, for example, Lipitor to the drug class of statins (HMG-CoA reductase inhibitors), within "lipid-modifying agents," in the cardiovascular category.[24]

III. Results

III.A. Sample Description

Summary statistics are shown in Table I. The overall sample is 90% white and 60% female. Roughly half is in fee-for-service Medicare (standalone PDP), with the rest in Medicare Advantage. As our sample is relatively young and not dual-eligible (excluding the poorest beneficiaries, for whom enrollment month does not affect cost sharing), mortality is low, 0.9 percentage points/year, and drug spending around $1,500/year. The 10 most used drug classes in our sample include statins, antihypertensives, diuretics, antidepressants, glucose-lowering drugs, and corticosteroids.

We check balance by regressing key pretreatment characteristics on enrollment month. Table II, columns (1) and (2) present means and estimated coefficients on enrollment month for the entire sample, respectively. Panel A shows that estimates for race, sex, and the number of drug prescriptions filled in the

23. We also find evidence that those born in January are less likely to delay enrollment and that those born in other months (e.g., November) are more likely to delay enrollment to January as opposed to other months.

24. We also attempted to measure medical diagnoses, procedures, and health care utilization besides drugs, using Medicare Parts A and B claims, including diagnoses, procedures, and admit/discharge dates. However, we are underpowered to detect effects for two reasons. First, the subsample of individuals enrolled in standalone PDPs (non-MA), for whom we observe these data, is half our sample. Second, 69% of all deaths happened outside of the hospital, implying a sudden event that did not result in prior utilization.

TABLE I

SAMPLE DESCRIPTIVE STATISTICS

	All	Spending group		
		Low	Middle	High
	(1)	(2)	(3)	(4)
Panel A: Demographics, spending, health				
White (%)	89.6 (30.5)	88.5 (31.9)	92 (27.1)	92.5 (26.4)
Female (%)	59.2 (49.1)	59.5 (49.1)	58.8 (49.2)	56.6 (49.6)
Standalone PDP (%)	51.7 (50)	46 (49.8)	64.2 (47.9)	71.2 (45.3)
Initial 90-day fills	5.35 (5.75)	3.16 (3.62)	9.75 (5.71)	16.9 (9.18)
One-year total spending ($)	1,478 (2,789)	626 (1,158)	2,828 (1,910)	9,185 (10,119)
One-year mortality (percentage points)	0.873 (9.3)	0.683 (8.24)	1.16 (10.7)	2.78 (16.4)
Panel B: Top 10 drug classes (%)				
Lipid modifiers	34 (47.4)	22.9 (42)	59.2 (49.1)	64.7 (47.8)
ACE inhibitors	20.5 (40.4)	16.4 (37.1)	30 (45.8)	30.8 (46.2)
Beta-blockers	18.9 (39.2)	13.3 (34)	31.4 (46.4)	37.7 (48.5)
Thiazide diuretics	18 (38.4)	14.4 (35.1)	26.7 (44.2)	25.7 (43.7)
Antidepressants	14.1 (34.8)	8.8 (28.4)	24.8 (43.2)	39.7 (48.9)
Corticosteroids	13 (33.6)	8.5 (27.9)	22.4 (41.7)	33.5 (47.2)
Acid blockers (GERD)	12.4 (33)	6.8 (25.1)	24.3 (42.9)	37 (48.3)
Anti-infectives	11.6 (32)	8.9 (28.4)	17.4 (37.9)	23.7 (42.5)
Hypoglycemics (oral)	11.2 (31.5)	6.4 (24.5)	20.9 (40.7)	33.2 (47.1)
Decongestants	11 (31.3)	6.7 (25)	20 (40)	31.6 (46.5)
Observations	358,706	251,093	96,849	10,764

Notes. Column (1) shows the mean (standard deviation) for the entire sample. Columns (2)–(4) show the same by initial 90-day spending group: low (1st–70th within-enrollment-month percentile), middle (71st–97th), high (98th–100th). One-year spending is measured from the first day of enrollment. One-year mortality is measured from December 1 of the first calendar year of enrollment to parallel our analysis. The percent on a drug class is measured by the presence of any claim in a given class in the first 90 days of enrollment. All participants are exactly 65 years old.

TABLE II

BALANCE OF KEY VARIABLES ACROSS ENROLLMENT MONTHS

	Entire sample		Middle spenders	
	Mean (1)	Enrollment month effect (std. err.) (2)	Mean (3)	Enrollment month effect (std. err.) (4)
Panel A: Demographics and key characteristics				
White (%)	89.6	0.00403	92	0.0219
		(0.0222)		(0.0378)
Female (%)	59.2	0.0582	58.8	0.165**
		(0.0356)		(0.0687)
Initial 90-day fills (count)	5.35	0.00364	9.75	0.00374
		(0.00419)		(0.008)
Predicted mortality (percentage points)	1.17	0.00119	1.38	0.00494
		(0.00139)		(0.00352)
Panel B: Initial mortality (cumulative percentage points, from enrollment)				
30 days from enrollment	0.058	−0.000053		
		(0.0017)		
60 days from enrollment	0.128	0.0013	0.082†	−0.0037
		(0.0026)		(0.0041)
90 days from enrollment	0.201	0.0037	0.181†	−0.0011
		(0.0032)		(0.006)

Notes. Panel A: Sample mean (column (1)) and coefficient on enrollment month (scalar), from regression of key pretreatment variables on enrollment month (an integer between 2 and 9; column (2)), for the entire sample ($n = 358{,}706$). Predicted mortality is estimated by fitting a model with demographics and initial (three-month) drug claims to predict subsequent (nine-month) mortality, in an independent sample of 66+-year-old dual enrollees. Columns (3) and (4) restrict to middle spenders ($n = 96{,}849$), based on spending in the first 90 days of enrollment. Panel B: Sample mean (column (1)) and coefficient on enrollment month (scalar), from regression of mortality in the first 30, 60, and 90 days after enrollment (column (2)). Columns (3) and (4) restrict to middle spenders, based on initial spending in the first 30 days of enrollment (unlike our main specification, which uses 90 days). Because we use the first 30 days to assign spending bins, we are unable to report the estimate for 0–30 days, and mortality rates in middle spenders reflect a one month shorter period than for the entire sample (i.e., 31–60 and 31–90 days, denoted by the †). Balance checks on initial spending are in Online Appendix Table A.2 and are discussed in Section III.A. * $p < .1$, ** $p < .05$, *** $p < .01$.

first 90 days are statistically and economically insignificant. Balance on spending is discussed above (Section II.C and shown in Online Appendix Table A.2). As a more synthetic test, we predict one-month mortality using all pretreatment variables and regress this on enrollment month.[25] This is also reassuring. Panel B shows a more direct balance check on mortality in the first three months of enrollment, by regressing mortality in the 30, 60, and 90 days after enrollment on enrollment month. Because cost sharing begins to affect enrollment months differently after this period (Online Appendix Table A.2), and the latest enrollment months has already entered December by the fourth month after enrollment, this is the last time we can compare mortality across enrollment months for the whole sample. No estimates are significant, providing further evidence that baseline health is similar between enrollment months. This provides some reassurance that conditioning our analytic sample on survival (i.e., to run our model of December mortality, beneficiaries must be alive on December 1) does not introduce selection bias correlated with enrollment month.

Having established balance, we now turn to defining initial spending bins as described above in Section II.C. Online Appendix Figure A.3 shows the effect of enrollment month on the likelihood of entering the donut hole by year-end, for each percentile of initial spending.[26] In the first 60 percentiles of initial spending, there is no relationship between enrollment month and donut-hole entry, and only a very slight negative relationship in the 61st–70th percentiles. Starting at the 71st percentile, we see significant and increasingly negative effects of enrollment month, as more and more earlier enrollees fall into the donut hole while later enrollees do not. Finally, at the 98th percentile, the rela-

25. This predictive model is estimated in a sample of 66+-year-old dual enrollees, who we assume have the same relationship between mortality and covariates. The independent variables are demographics (race, sex) along with the consumption metric measured over the first 90 days of coverage for the year (January–March), to mirror our main sample, and the dependent variable is mortality (April–December). We apply this model to generate predictions in our main sample, using covariates measured over the first 90 days of enrollment.

26. As discussed in Section II.C and Online Appendix Table A.2, we use 90-day initial spending because it is balanced across enrollment months, while spending past 90 days is likely affected by forward-looking behavior. There are not unique percentiles for the 70% of initial spending so we estimate by decile for this group. We find precise null estimates for all six deciles.

tionship begins to reverse, and the last percentile bin has a large positive effect, as earlier enrollees exit the donut hole and enter the catastrophic coverage.

Our primary goal in setting spending-group cutoffs is to estimate equation (4) in a homogeneous group of middle spenders, whose drug consumption monotonically increases in enrollment month. Estimates for low spenders and high spenders are not our primary focus, so we are willing to tolerate some heterogeneity in these groups to ensure homogeneity in the middle-spender group. We thus assign the first 70 percentiles to the low-spender bin—even though empirically this is likely to include a small number of lower middle spenders from percentiles 61–70 who are slightly affected by the donut hole—and percentiles 98–100 to the high-spender bin. On average, the 98th–99th percentiles likely experience the impact of enrollment month more like a middle spender than a high spender (i.e., the effect of enrollment month on donut-hole entry is negative); because our primary interest is in estimates from the middle-spender bin, we set the cutoff very conservatively to minimize non-monotonicity of enrollment month in that bin. We performed a sensitivity analysis over a range of alternative cutoffs (e.g., starting the middle-spender bin at the 61st percentile or ending it at the 95th percentile) and found nearly identical results, shown in Online Appendix C.1.

Table I, columns (2)–(4), present summary statistics by spending group. Middle spenders, our main population of interest, are more likely to be fee-for-service and have higher drug spending (by construction) and mortality than the overall sample mean. Table II, columns (3) and (4), present balance checks within the middle-spending group. Although we do find a statistically significant association between enrollment month and sex in the middle-spenders subsample, it is quite small in percentage terms (0.29%) and orders of magnitude too small to explain the mortality differences we find. Drug consumption, measured by the number of prescriptions filled in the first 90 days, is balanced across enrollment months. Our most important balance tests—predicted mortality based on observables and actual initial mortality in the first months of enrollment—are likewise reassuring.

III.B. Mortality Effects of Drug Interruptions

Figure III shows our identification strategy and main result graphically, using a set of indicators for enrollment month. In

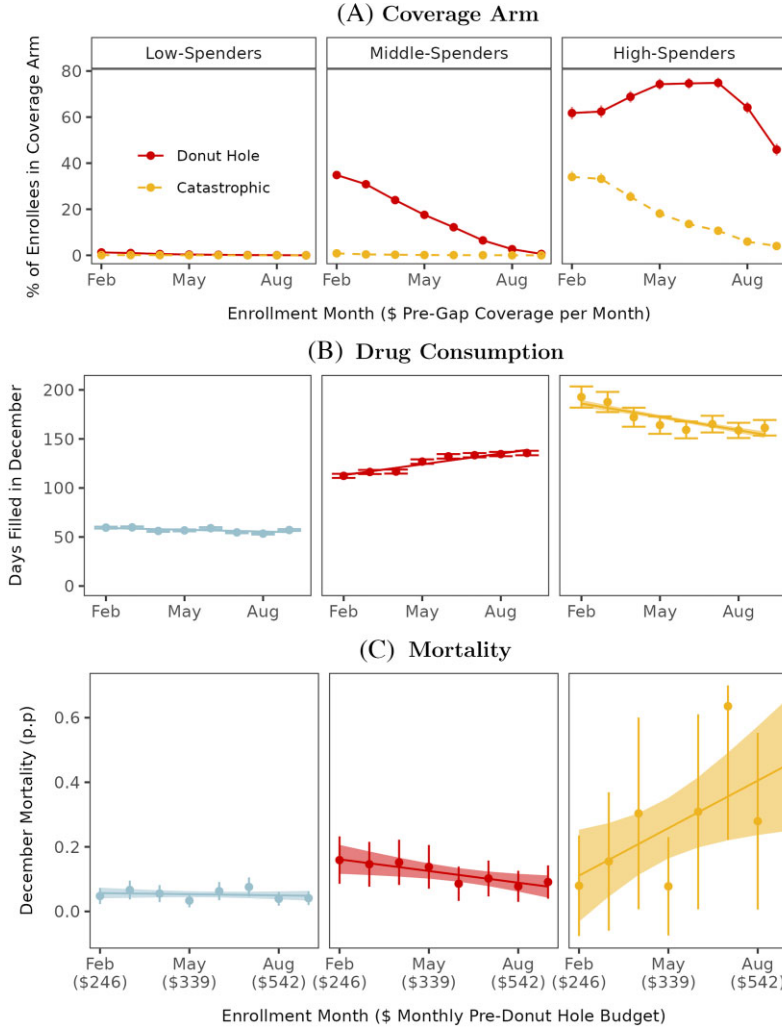

FIGURE III

Effect of Enrollment Month on End-of-Year Coverage Arm, Drug Consumption, and Mortality

Panel A: Fraction in the donut hole or catastrophic coverage at year-end (y-axis) by enrollment month (x-axis) and initial 90-day spending bin (horizontal subpanels). Panel B: Drug-days filled in December (y-axis) by enrollment month and spending bin. Panel C: December mortality (y-axis) by enrollment month (translated into monthly pre-donut budget, in parentheses) and spending bin. Ninety-five percent confidence intervals are shown; one (July, high spenders) is truncated.

Figure III, Panel A, we see the proportion of beneficiaries in each coverage arm in December, by spending group. This aggregates the finer percentile bins in Online Appendix Figure A.3 and shows a similar pattern of clear separation of (mostly) low spenders on the left, for whom enrollment month has no effect on donut-hole entry; middle spenders, where early enrollees end up in the donut hole more often (34.9% of February versus 0.6% of September enrollees) but not the catastrophic coverage; and (mostly) high spenders on the right, where early enrollees first enter the donut hole more often, then also enter the catastrophic coverage more often (the latter cannibalizes the share of early enrollees in the donut hole, hence the non-monotonic trend for the former).

Figure III, Panel B shows that these differences in donut-hole exposure affect drug consumption, as measured by December drug-days filled (i.e., the number days supplied, summed across all drugs). As expected, low spenders have no variation in days filled by enrollment month. Middle spenders fill substantially fewer days when they enroll early: February enrollees, for example, fill 25.5 fewer drug-days in December than September enrollees. High spenders show the opposite pattern, with February enrollees filling 31.8 more drug-days in December than September enrollees.[27] We take all estimates in this latter group with a grain of salt, due to the non-monotonicity of the enrollment-month effect on consumption; the complexity of interpreting these coefficients means these beneficiaries are not the focus of our analysis. Online Appendix Table D.1 shows the specific drug classes most affected by price increases, which include statins, beta-blockers, ACE inhibitors, and antidepressants; notably, opioids are unaffected.

Figure III, Panel C shows the relationship between December mortality and enrollment month. For low spenders, there is no effect. For middle spenders, we find a significant and large negative relationship between mortality and enrollment month that mirrors the increase in drug consumption shown in Panel B. For high spenders, we see a large positive relationship, again contrasting with middle spenders. Table III, Panel A summarizes these figures, and reports OLS estimates of β_2 from equation (4), both in terms of enrollment month (column (2)) and monthly pre-

27. Figure IV shows that these decreases are indeed interruptions, particularly at year-end: differences across enrollment months are small earlier in the year, then accelerate dramatically, peaking in December.

TABLE III

MORTALITY EFFECTS OF PRICE-DRIVEN DRUG INTERRUPTIONS

	December mortality (percentage points) (1)	Enrollment-month effect (percentage points/month) (2)	Pre-donut budget effect (percentage points/$100) (3)
Panel A: Enrollment month effect, by initial spending			
Low spenders	0.053	−0.000362	−0.000771
		(0.00199)	(0.00317)
Middle spenders	0.118	−0.0112**	−0.0164**
		(0.00498)	(0.00786)
High spenders	0.288	0.0469**	0.0549
		(0.0219)	(0.0402)
Difference: middle vs. low	—	−0.0109**	−0.0156*
		(0.00535)	(0.00846)
Difference: middle vs. high	—	−0.0581***	−0.0713*
		(0.0225)	(0.0411)
Panel B: Pre- vs. post–donut hole closing (middle spenders)			
Full donut hole (2007–10)	0.120	−0.0137**	−0.021**
		(0.00626)	(0.00993)
Closing donut hole (2011–12)	0.113	−0.0063	−0.00993
		(0.00814)	(0.0129)
Difference: full vs. closing	—	−0.00745	−0.0111
		(0.0103)	(0.0163)

TABLE III
CONTINUED

	December mortality (percentage points) (1)	Enrollment-month effect (percentage points/month) (2)	Pre-donut budget effect (percentage points/$100) (3)
Panel C: Key falsification estimates (middle spenders)			
Main Sample: December, age 66	0.12	0.00413	0.0075
		(0.00496)	(0.00826)
Duals: December, age 65	0.38	−0.0013	−0.0144
		(0.0132)	(0.0205)
Disabled: December, age 64	0.57	0.0045	−0.00661
		(0.0211)	(0.0316)

Notes. Panel A: December mortality rate (column (1)) and coefficient γ_2 on enrollment month from equation (4) (column (2)), by initial spending: low spenders (lowest 70% of initial 90-day spending: unlikely to enter the donut hole irrespective of enrollment month); middle spenders (71st–97th percentile: more likely to enter the donut hole if enrolling earlier; and high spenders (98th–100th percentile: likely to enter both the donut hole and then the catastrophic coverage if enrolling earlier). Column (3) translates enrollment month into pre-donut budget (in $100/month) before full cost sharing. Panel B: Coefficient on enrollment month and pre-donut budget for middle spenders ($n = 96{,}849$), before versus after the donut hole began to close in 2011. Difference rows are pairwise tests for equality of coefficients. Panel C: Selected falsification tests, estimating enrollment-month effect on mortality in settings with no policy link between enrollment month and drug prices. Robust standard errors are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$.

donut budget (column (3)) as scalars. Concretely, among middle spenders, a $100 budget increase—24.4% relative to the average enrollment-month budget in our sample—leads to a mortality reduction of 0.0164 percentage points, or 13.9% of the base mortality rate. In low spenders, the effect is small and insignificant. In high spenders, the effect is positive, large, and significant.[28]

A central question is whether variation in mortality across enrollment (or birth) months is confounded by factors other than cost sharing. Recall that we have already seen two reassuring facts in this regard. First, no such variation in mortality exists in the first few months after enrollment, before cost sharing and drug consumption trends start to diverge across enrollment months (see Table II, Panel B). Second, any confounder would have to correlate not only to enrollment month, but also to the exact combination of enrollment month and initial spending patterns we would expect, given Medicare policy: large for medium spenders, absent for low spenders, and opposite sign for high spenders.

Nevertheless, we directly address two known potential confounders. The first and most obvious is age: those who enroll earlier in the year are older than those who enroll later, so our results indicate older enrollees die more—not a novel finding. We emphasize, however, that if aging alone were responsible, this trend would not be affected by initial spending differences. In addition, a simple calculation using U.S. life tables (from the Social Security Administration) illustrates why this is unlikely to be a concern. Comparing annual mortality rates for 65- versus 66-year-olds, we estimate the effect of being one month younger is to affect monthly mortality by roughly −0.001 percentage points, or −0.76%.[29] This is quite similar to the effect of enrollment month from the low spenders, albeit imprecisely estimated, which in the absence of enrollment-month effects on budget are likely to be

28. The significant coefficient is in the enrollment-month specification, but not the pre-donut budget specification, which is not surprising. We would not expect the latter to be a good way to scale enrollment month for those who have greatly exceeded the donut-hole budget and entered the catastrophic phase.

29. Using 2010 data, annual mortality was 1.59 percentage points for 65-year-olds and 1.74 percentage points for 66-year-olds, translating into monthly rates of roughly 0.133 percentage points and 0.146 percentage points. We interpret the difference, −0.013 percentage points, as the effect of being one year younger on monthly mortality. We translate this annual effect into a monthly effect, then divide by mean monthly mortality to get the percent decrease.

solely attributable to age: −0.00036, or −0.68% (Table III). Another data point comes from a null distribution of effects of enrollment (or birth) months on mortality, across many observably similar samples lacking enrollment-month effect on drug budgets (which we describe in Section III.C). This yields a median estimate of −0.53%, which again likely reflects the effect of being one month younger in these populations. In other words, several different methods of calculating the aging effect all give fairly consistent estimates, between −0.5 and 1%—an order of magnitude smaller than the effect size from our main analysis, which captures the effect of enrollment month on cost sharing plus the effect of being one month younger: −0.118 percentage points or −9.49%. (The equivalent estimate based on birth month is −0.009 percentage points, or −7.7%, shown in Online Appendix Table C.2, Panel A.) This gives us confidence that aging is only a small part of the relationship we observe.

A second potential confounder is health differences by birth season, which has been suggested to result from disease seasonality (Currie and Schwandt 2013) or selection (Buckles and Hungerman 2013).[30] Most of the literature focuses on peri- and postnatal outcomes, but two large studies explore later-life outcomes in populations similar to our own. In a study of life expectancy at age 50, Doblhammer and Vaupel (2001) find that mortality peaks among May births, a pattern that is consistent across multiple cohorts in the Northern Hemisphere (Austria and Denmark; it is exactly reversed in Australia). This echoes a Nurses' Health Study by Zhang et al. (2019) that finds cardiovascular disease mortality peaks among April births (although that study found no overall mortality differences). In our sample, by contrast, mortality peaks among February births, and May births are in the middle of our distribution of mortality effects (Online Appendix Table C.2, Panel B shows these effects by birth month, rather than enrollment month, for comparability to this literature). In addition, our mortality effects are again orders of magnitude larger.[31] To summarize, the birth seasonality litera-

30. While a full review of this literature is beyond the scope of our work, we refer the reader to Currie and Schwandt (2013) for an excellent summary and a very rigorous empirical exploration of mechanisms.

31. Doblhammer and Vaupel (2001) find the maximum effect of one birth month change on life expectancy (at age 50: e_{50}) is 0.05–0.1 year, versus an average remaining e_{50} of 27.5 years. Without access to the full life table, converting

ture demonstrates small, cyclical increases in mortality among those born in April and May, while we find large mortality increases earlier in the year, that increase linearly with enrollment month, and that are larger and correlate with Medicare drug budgets.

III.C. Falsification Tests

As shown in Figure III, the effect of enrollment month on mortality varies across spending bins, following the pattern of idiosyncratic changes in drug budgets created by Medicare policy. We consider this our first falsification check because it would be so hard for potential confounders to match the very specific set of idiosyncratic budget limits, initial spending, and calendar-month effects we exploit. We build on this insight to develop a more comprehensive set of falsification checks that replicate our analysis in a range of closely related settings lacking an enrollment month (or birth month)–drug budget link. This lets us further verify that our findings are only seen in the presence of Medicare budget limits on drug consumption.

First, Figure IV isolates the mortality effect to the end of the first calendar year: it appears just as differences in drug consumption peak in December and disappears as soon as prices reset and drug consumption reequalizes in January. The figure is constructed by following our main population of middle spenders over time, before and after December. In each calendar month, we summarize the effect of enrollment month on being in the donut hole (Panel A), drug consumption (Panel B), and mortality (Panel C) via a linear coefficient. For example, the red point in Panel A is the coefficient on enrollment month, from a regression of donut-hole entry in December on enrollment month. This is analogous to ϕ_1 from equation (3), but estimated in all middle spenders. In other words, this red point summarizes the middle panel of Figure III, Panel A as a single linear coefficient. Each point on the graph similarly represents an estimate of the enrollment-month effect on an outcome, one for each calendar month, from August

between e_{50} and annual mortality risk is not possible, but using the rule of thumb from Pollard (2002), these e_{50} changes imply on the order of a 0.03% relative annual change in mortality risk (in each one-year age bin), compared with the effect of 10.4% we observe (0.013 percentage points per birth month versus base mortality of 0.124 percentage points for middle spenders, Online Appendix Table C.2, Panel B).

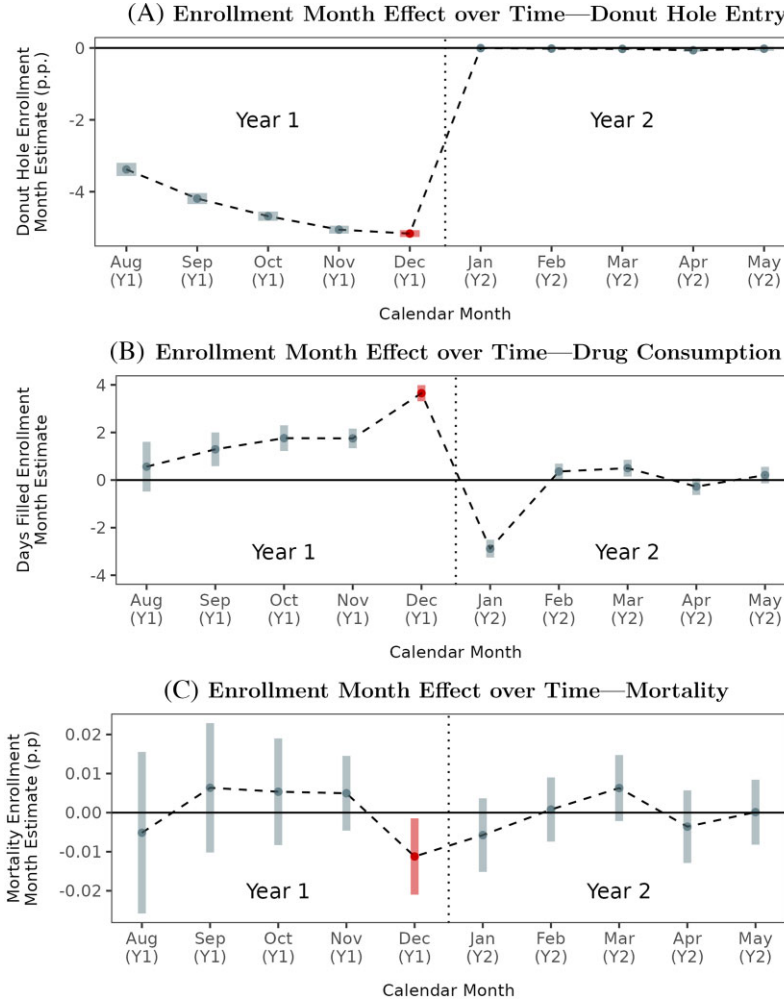

FIGURE IV

Evolution of Enrollment Month Effects over Time, August (Year 1)–May (Year 2)

Each point shows the linear coefficient (and 95% confidence interval) measuring the effect of enrollment month over time, from August of calendar year 1 to May of year 2, for middle spenders. Panel A shows the effect on donut-hole entry, Panel B drug consumption, and Panel C mortality. The red points show enrollment month effects in the calendar month of December (i.e., each point summarizes one of the three middle panels of Figure III, as a linear coefficient). For example, Panel C shows γ_{1t} from the regression $Y_{it} = \gamma_{0t} + \gamma_{1t}Z_i + X_i\gamma_{2t} + \gamma_{year} + \gamma_{plan} + \epsilon_{it}$, for middle spenders in months $t = \{8, 9,..., 12, 1, 2,..., 5\}$, and the red point is $\gamma_{1,12}$.

of the first calendar year to May of the second year.[32] Panel A shows that the effect of enrollment month on donut-hole entry grows smoothly over time in year one, as more beneficiaries enter the donut hole. It then disappears in January of year two, when all enrollment months reenter the initial coverage phase.

Similarly, each point in Panel B is the estimated linear coefficient of monthly drug-days filled on enrollment month (analogous to α_1 from equation (2), but estimated in all middle spenders). Differences in consumption between early and late enrollees appear as early as September, and this effect grows steadily over the next three months. But in contrast to Panel A, the magnitude abruptly jumps up in December: the effect on drug consumption more than doubles (1.75 drug-days difference per enrollment month in November, versus 3.65 in December). These consumption gradients completely reverse in January: as soon as prices reset, earlier enrollees—who have been waiting out the high prices of the donut hole and filling fewer drug-days December as a result—now take advantage of lower prices to make up for their missed doses. Indeed, not only do earlier enrollees fill 2.86 more drug-days per enrollment month, they also fill sooner (0.30 days sooner for each earlier enrollment month, conditional on filling in January; estimate not shown). These patterns fit with the beliefs documented in our survey (Figure I): patients view short interruptions in their drugs as largely innocuous, and put off filling until prices reset.

Finally, Panel C presents estimated linear coefficients of monthly mortality on enrollment month ($\gamma_{1\mathcal{M}}$ from equation (4)). Recall that our balance checks showed no differences in mortality in the first three months of enrollment. This panel shows a similar analysis, by calendar month rather than enrollment month. We find no significant mortality gradient across enrollment months from August to November—then a large significant effect in December. In January, just as prices reset and earlier enrollees rush to fill their medications, the mortality effect attenuates: it is negative but insignificant in January, then disappears altogether from February onward. Overall, these results show that mortality increases are a transient phenomenon tied

32. For months prior to December, we exclude those who have not been enrolled for three months, to identify a similar group of middle spenders over time (e.g., the September regression includes February–June enrollees). We could not produce stable estimates prior to August of year one, as the sample size decreases with each month due to fewer enrollment months.

to abrupt increases in drug interruptions in December, and December alone. This is close to what we would expect based on our effect size and power calculations (Online Appendix Figure C.3).

Our second falsification test leverages policy variation over time, to show that the mortality effect is proportional to the degree of cost sharing mandated by evolving Medicare policy. In 2011, the donut hole began to close, allowing us to compare the mortality effects of enrollment month in our main sample of middle spenders, before versus after a policy change that reduced cost sharing faced by earlier enrollees. Table III, Panel B shows estimates based on equation (4), but with enrollment month interacted with indicators for pre– versus post–policy change and restricted to middle spenders. The effects are larger before the attenuation of the donut hole, with a mortality reduction of 0.0210 percentage points per $100/month pre-gap budget increase, isolated to 2007–2010. In other words, the effect of cost sharing on mortality is concentrated in the enrollees we expect, and over the time period when these enrollees are most affected by cost sharing.

Our most comprehensive set of falsification tests situates our main estimate in an empirical "null distribution" of enrollment (or birth) month effects on mortality, from a range of related populations and time periods. The placebo effects are estimated in samples of Medicare beneficiaries who are similar to those in our estimation sample, but lack the idiosyncratic link between enrollment month and drug budgets. As a result, we expect this distribution to be centered at zero (adjusting for the effect of age, which we discuss in detail below), and for our main estimate to be in the extreme left tail. Our first set of estimates extends the analysis in Figure IV to follow our main sample further in time, estimating monthly effects from January of their second calendar year of enrollment until December of their fourth year (36 estimates: shown separately in Online Appendix Figure C.1, Panel A). Our second set replicates the analysis in older dual eligibles, who do not face cost sharing. We pool all years together, then split into subsets defined by demographic factors and geography (49 estimates: Online Appendix Figure C.1, Panel B). Third, we broaden to a larger set of beneficiaries 66 years old and above—non-duals, older duals, and disabled dual beneficiaries (ages 50–64), none of whom face the exact same cost sharing as 65-year-old non-duals—whose initial spending makes them observably similar to our mid-

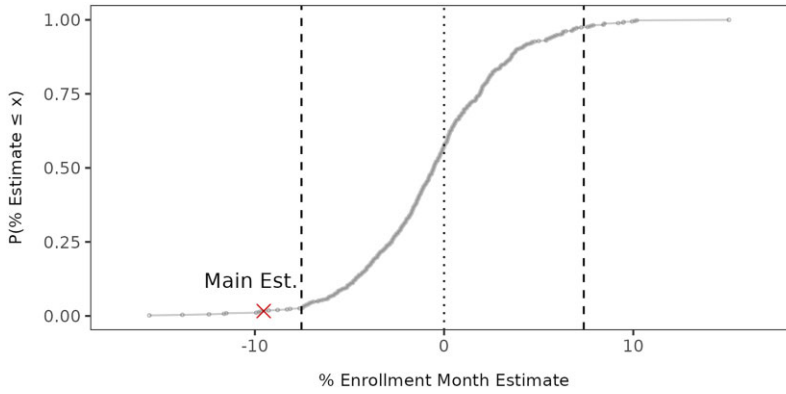

FIGURE V

Distribution of "Placebo" Estimates of Enrollment or Birth Month on Mortality

Regression estimates of the effect of enrollment month or birth month on mortality, for middle spenders, in a variety of settings lacking an enrollment month–drug budget link: non-dual enrollees age 66–85, dual enrollees from 66–85, and disabled enrollees age 50–64. Estimates are divided by mean mortality in each sample to get a percentage change. This figure pools all falsification tests together; Online Appendix Figure C.1 provides further detail on the separate types of tests that contribute. Vertical lines show the 2.5th and 97.5th percentiles. Our main (nonplacebo) estimate from Table III is shown as a red cross.

dle spenders (459 estimates: Online Appendix Figure C.1, Panel C).[33]

Figure V reports the distribution of results across these 544 falsification samples. We rank estimates of the enrollment month effect by magnitude on the x-axis.[34] The y-axis shows the cumulative fraction at least as large as x. The median estimate is -0.54%, likely reflecting the effect of age across these samples: it is quite close to the age effect we estimated from other sources in Section III.B (-0.76% from Social Security data, or -0.68% from the low spenders). The estimate from our main analysis, -9.49%, is shown as a red cross. It is at the 97.8th percentile of mortality effects, larger in absolute magnitude than 532 of 544 placebo

33. Because some of these populations lack an observable enrollment month, we use birth month as a proxy. More details are in the Online Appendix.

34. Estimates are relative to the baseline mortality in each sample for comparability: falsification samples vary in their baseline mortality. Many samples are sicker than our 65-year-old non-dual primary sample, because of older age, lower income, and enrollment based on disability. As a result, mortality is higher, which if anything could make us better powered to detect a (spurious) effect.

TABLE IV
DRUG CONSUMPTION CHANGES, BY PREDICTED RISK

	All		Lowest 2/3 risk		Top 1/3 risk	
	Mean (1)	Est. (std. err.) (2)	Mean (3)	Est. (std. err.) (4)	Mean (5)	Est. (std. err.) (6)
All classes	126.40	4.93*** (0.289)	110.70	3.43*** (0.317)	156.80	7.78*** (0.559)
Cardiovascular	50.20	1.42*** (0.157)	41.80	0.766*** (0.174)	66.50	2.65*** (0.304)
Diabetes	10.20	0.618*** (0.0701)	10.10	0.552*** (0.0862)	10.40	0.746*** (0.12)
Respiratory	5.30	0.459*** (0.0453)	5.10	0.421*** (0.0548)	5.60	0.53*** (0.0793)

Notes. Effect of cost sharing on December drug consumption for key drug classes. Cardiovascular drugs includes statins, beta-blockers, ACE inhibitors, calcium channel blockers, angiotensin receptor blockers, and thiazide diuretics. Diabetes includes both insulin and oral hypoglycemic agents. Respiratory includes inhaled and oral treatments for chronic pulmonary disease. Column (1) shows mean number of drug-days filled by middle spenders in December (days supply, summed across all prescriptions filled, and grouped by drug class). Column (2) presents estimates (and robust standard errors) of the effect of pre-donut budget (in $100s) on drug-days. Columns (3)–(6) show mean drug days and similar regression estimates, by risk of the adverse events each drug class prevents: heart attack and stroke for cardiovascular drugs, diabetic complications for diabetes, and respiratory failure for pulmonary drugs. (For "all classes," we use predicted cardiovascular event risk, since both cardiovascular drugs and cardiovascular mortality are the most common.) * $p < .1$, ** $p < .05$, *** $p < .01$.

estimates overall (and more negative than 536 of 544 estimates, or 98.5%). Online Appendix Figure C.2 shows similar results in a plot of t-statistics from these regressions. This omnibus test complements the more tailored falsification tests above and builds confidence that our observed mortality effect is in fact due to difference in the pre-donut budget faced by enrollees, rather than any spurious correlation.

III.D. Are These Effect Sizes Medically Plausible?

Table IV shows that a consistent finding of prior work is true in our setting as well (Brot-Goldberg et al. 2017; Einav, Finkelstein, and Polyakova 2018). Patients cut back on many drugs that have been shown to prevent life-threatening adverse events in randomized trials: cardiovascular (e.g., statins, antihypertensives), glucose-lowering (e.g., insulin), and respiratory (e.g., steroids, inhalers). For these classes, column (1) shows the fraction of middle spenders who ever fill the drug. Column (2) shows estimates of how a $100 increase in monthly pre-donut budget affects drug-days consumed in December. Earlier enrollees fill 5.2

more drug-days overall, with half of this total accounted for by cardiovascular, diabetes, and respiratory drugs (1.5, 0.6, and 0.5 more drug-days, respectively).[35]

Could short interruptions in these drugs cause mortality effects of the magnitude we find? Two sets of facts in the medical literature indicate that this is plausible. First, it is a common misperception that drugs for chronic diseases work slowly. Certainly, clinical trials for these drugs last many years—but we should not conflate the time scale required to measure effects for rare, noisy outcomes like mortality with the time scale on which effects begin. Inspection of many published survival curves for chronic drugs shows that they start to diverge almost immediately but only reach statistical significance after years. Statins provide an instructive example: the Kaplan-Meier curves in the landmark JUPITER trial begin to diverge at the origin (Ridker et al. 2008) and, as Heeschen et al. (2002) note, appear to have short-term protective effects in patients during hospitalizations for heart attack. This seems counterintuitive, since the cholesterol-lowering effect of statins is to reduce atherosclerosis (heart disease) over long time periods. But statins also act via a range of other pleiotropic mechanisms: they prevent blood clotting and reduce inflammation and reactivity of blood vessels (Oesterle, Laufs, and Liao 2017). These mechanisms yield large protective effects in the very short term for patients with acute conditions like heart attack and stroke, which is why clinical guidelines mandate initiating a statin immediately on diagnosis of heart attack as a result (Rosenson 2023). Several other chronic medications have similar effects over multiple time scales: the diabetes drug metformin lowers blood glucose in the short term and has a variety of longer-term anti-aging effects (Kulkarni, Gubbi, and Barzilai 2020); antibiotics used for chronic obstructive pulmonary disease treat acute infections and reduce long-term inflammation (Blasi, Mantero, and Aliberti 2012). As a result, even short drug interruptions can mean large forgone treatment bene-

35. We caution against any effort to tie the magnitude of changes in estimated consumption to changes in estimated mortality in the same time period, given the bias discussed in Section III.B: lagging mortality effects of drugs, combined with intertemporal substitution across periods, means that scaling the mortality effect by consumption will bias estimation of the effect of consumption (and thus consumption changes due to cutbacks).

fits, depending on the idiosyncratic time scale of the mechanisms mediating the drug's treatment benefit.

Second, inferring the effect of drug interruptions from the effect of drug initiation may be misleading in the presence of rebound effects. Drugs induce a complex set of physiological changes that put patients in a new physiological equilibrium—indeed, that is the point of taking drugs. Abruptly stopping a long-standing drug, to which the body has adapted, can precipitate a potentially dangerous set of effects in several settings. This idea has entered the popular consciousness in the setting of opiate withdrawal, but opiates are far from the only drug class with such effects. Indeed, rebound effects have been noted for 7 of the 10 most commonly taken drugs in our sample (Table I)—statins (Heeschen et al. 2002), antihypertensives (Psaty et al. 1990), diuretics (Walma et al. 1997), antidepressants (Horowitz et al. 2021), corticosteroids (Jarad et al. 1999), and glucose-lowering drugs (Czosnowski et al. 2009). Because of the practical and ethical difficulties of studying drug interruptions directly, the best evidence for guidelines on drug tapering—which recommend slow transitions under close medical supervision (Bain et al. 2008; Steinman and Reeve 2023)—comes largely from in vitro experiments, or older, idiosyncratic studies.[36] Our results add new weight to these recommendations, suggesting that drug interruptions can precipitate serious adverse events. Interestingly, we find that 69% of December deaths in our sample occur outside the hospital, suggesting a catastrophic, sudden event.[37]

If our results fit with medical intuition, they are more discordant with some parts of economics, as they run counter to the predictions of several standard models of behavior. Consider a population of patients prescribed a drug by a doctor, for whom treatment benefit is heterogeneous across individuals, as research indicates (Chandra and Skinner 2012; Chandra and Staiger 2020). If the price of that drug suddenly increases, traditional models predict that cutbacks should be concentrated among people for

36. Even for beta-blockers, perhaps the class of drugs for which there is the most evidence, a recent article can be summed up by its title: "Beta blocker rebound phenomenon is important, but we do not know its definition, incidence or optimal prevention strategies" (Koracevic et al. 2020).

37. This estimate comes from enrollees in standalone PDPs, 52% of our sample, in whom we also observe hospitalizations. This combination of high Medicare Advantage prevalence and high out-of-hospital death rate means we are underpowered to detect changes in hospitalizations in this study.

whom the drug yields fewer benefits. Under moral hazard, the marginal drugs that patients drop are disproportionately low benefit. Likewise, a Roy model of patient decision making with private information on heterogeneous treatment effects predicts that patients self-select into treatments that benefit them more. The corollary of this is that those with the highest potential benefit should be willing to pay more for a drug and should thus cut back less when the price increases. So a key question here is: who is interrupting their consumption? If interruptions occur largely in patients with low treatment benefit, just as moral hazard or patient private information would predict, it would be hard to square with our results. If, on the other hand, high-benefit patients interrupt their consumption due to relatively small price increases, it would help make sense of the large mortality effects we see; it would also raise a fascinating new set of questions about why this might happen.

To develop a measure of an individual's health benefit from a drug, we use machine learning to form predictions on patient risk. Focusing on a set of drugs used to prevent key adverse events, we assume that the benefit of a given drug is proportional to the baseline risk of those outcomes that the drug prevents. For example, we assume the benefit of a statin is proportional to the risk of heart attack and stroke. This assumption is supported by both a substantial body of evidence and specific clinical guidelines, particularly for cardiovascular drugs. Major randomized trials (e.g., JUPITER, HOPE-3, CARDS, and ASCOT, reviewed by Bibbins-Domingo et al. 2016) show 30%–50% larger absolute risk reductions from statins in groups with higher predicted risk of heart disease, whether defined by age, diagnosed risk factors (e.g., diabetes), or biomarkers (e.g., LDL, CRP). Studies of polygenic risk scores show similar heterogeneity, with higher-risk participants getting nearly three times the absolute risk reduction (Natarajan et al. 2017). Clinical guidelines also reflect this assumption, for example, the American College of Cardiology 10-year risk calculator to guide treatment for cardiovascular disease. There is similar medical consensus and biological plausibility, if less strong empirical evidence, for diabetes and respiratory drugs.[38]

38. Even if this model is far from optimal, in the sense that it captures true treatment heterogeneity, if patients or doctors believe that high risk equates to high benefit, this measure will identify patients who believe they would benefit from a given treatment. We view this as a useful fact to understand.

Concretely, we identify three important drug classes M—cardiovascular, diabetes, respiratory—and compile a list of observable adverse outcomes that the drugs are prescribed to prevent: heart attack and stroke for cardiovascular medicines, diabetic complications (e.g., foot amputation) for hypoglycemic medicines, and respiratory failure for inhalers and steroids. This allows us to define indicators Y_M, one for each drug class, that indexes whether a beneficiary experienced an adverse event preventable by drug M. We form separate predictive models for each outcome and restrict to those who are not taking class M (e.g., when predicting risk of heart attack or stroke, we exclude patients on statins and other medications for coronary artery disease), to obtain a prediction of the risk of complications if untreated.[39] These models use a beneficiary's demographics and initial 90-day claims to predict the likelihood of adverse events over the next 270 days, and are trained on an entirely separate sample of dual-eligible 66+-year-olds, to ensure our predictions are out of sample. We turn the model's continuous risk predictions into simple indicators, $\hat{Y}_{Mi}$, that index the highest-risk one-third of the sample, based on where risk begins to increase rapidly (see Online Appendix Figure D.1). Additional details are in Online Appendix D.

Strikingly, the highest-risk beneficiaries cut back at least as much, if not more, on those medications that benefit them the most. We run equation (2) with enrollment month interacted with $\hat{Y}_{Mi}$, separately for each drug class M and restricting to middle spenders as usual. Table IV shows the results, which are especially pronounced for cardiovascular drugs: for each $100/month budget decrease, low-risk patients fill 0.766 fewer cardiovascular drug-days (a 1.8% reduction), while high-risk patients fill 2.65 fewer (4.0%).[40] This finding is incompatible with standard eco-

39. In potential-outcomes notation, we wish to predict Y_M^0, not Y_M^1, as our proxy for the benefit of drug M. Naturally this choice of prediction target also induces selection bias, as noted in Mullainathan and Obermeyer (2019), who use machine learning to predict the yield of testing for heart attack in the tested and then validate the model in the untested. Building on that work, Online Appendix Figure D.1 shows that true risk rises monotonically in predicted risk for the treated just as the untreated, establishing face validity of the predictor irrespective of treatment status.

40. We emphasize the relative, not the absolute, magnitude of these changes, which reflect only December differences in consumption and not cumulative differences over the year; see Online Appendix B.

nomic models of behavior grounded in moral hazard (Pauly 1968; Einav and Finkelstein 2018) or private information (Finkelstein and McGarry 2006; Hendren 2013): those at high risk of a cardiovascular event should have the highest and most inelastic demand for treatment, reflecting their benefit from the drug.[41] We find similar but less pronounced trends for diabetes and respiratory drugs.

Importantly, Online Appendix Table D.2 demonstrates that these effects are of similar magnitude in high- and low-income ZIP codes alike. So while socioeconomically disadvantaged patients may have both worse health and less ability to pay for treatments, our results are unlikely to be driven by socioeconomic factors alone.[42] We do identify one potential contributor to this effect: a subgroup of beneficiaries chooses to fill no drugs when prices increase—no matter how many drugs they were on prior to the price shock or their individual health risks. Indeed, those with higher initial consumption are in fact more likely to interrupt their entire regimen of prescribed drugs.[43] Online Appendix Table D.3 shows that enrollment month makes it more likely for all middle spenders to fill no medications (Panel A), but also shows that those in the higher two terciles of initial drug consumption (in the first 90 days of enrollment) are far more affected in absolute and relative terms. Mechanically, this behavior results in large absolute reductions in drug use for higher-risk patients, who are on more drugs to begin with.

41. As noted already, our population of middle-spending patients have been prescribed a basket of drugs and then must decide whether to continue them after a sudden price increase. Among this population, which is already filling 126 drug-days per month, lower- versus higher-risk patients should be more elastic. These results do not apply to low spenders, who are at such low risk that they are not prescribed any medications, and who would also have zero demand for treatment regardless of price.

42. Of course, there is variation in income within ZIP codes, often quite a bit, so this does not rule out income effects or liquidity constraints. However, to the extent that we see similar behaviors in rich and poor areas alike, it forms some upper bound on how important these effects can be on average.

43. This is reassuring that this phenomenon is not simply a floor effect, that is, due to left censoring: it is more common, not less, in those with more drug fills to begin with.

IV. Conclusion: Errors and Misinformation

Clinical medicine and economics share a fundamental respect for individuals' preferences and decision making. Our results, alongside a substantial literature in both fields, present a dilemma for this perspective, if individuals make decisions that are inconsistent with the general preference to stay alive. A rough calculation highlights the incongruous life-year valuations implicit in the price-driven cutbacks we study. We divide the effect of enrollment month on cumulative drug spending (from September to December, when drug-day differences begin to emerge across enrollment months), by the effect of enrollment month on mortality to infer the life-year valuations underlying the decision to interrupt drugs in our sample: $11,321 (95% confidence interval: $6,195–$73,858).[44] Put another way, at a widely used life-year valuation of $100,000 a year (Neumann, Cohen, and Weinstein 2014), a 65-year-old middle spender in our sample would have to believe that she had at most 2.19 years left to live. This contrasts sharply with average life expectancy in the general population at age 65—19.2 years—and observed outcomes in our sample: at a median follow-up period of five years, 93.2% of middle spenders are still alive. This is hard to square with the idea that patients are equalizing marginal benefit with marginal cost of drugs, particularly in the absence of obvious ZIP code differences that might indicate income or liquidity constraints. It supports the idea that the price elasticity of demand is an insufficient statistic for welfare, as has been noted by Baicker, Mullainathan, and Schwartzstein (2015) and Einav and Finkelstein (2018). This "behavioral hazard" has far-reaching implications for the design of health insurance, particularly as insurers place more emphasis on cost sharing.

Behavioral economics provides several potential explanations for our results, in the form of predictable distortions in the

44. This is based on the full cost of the drug, which is an upper bound on how much patients actually pay. We estimate a 2SLS regression of December mortality on instrumented spending from October–December of year 1, in middle spenders (we emphasize that this is a very approximate exercise and all the caveats from Online Appendix B apply). The inverse of this estimate is dollars per life-year,which we divide by average life expectancy at 65 from Social Security data (weighted by the proportion of males/females in our sample) to estimate life-year valuation. Finally we calculate the implied life-year valuation at the bounds of the 95% confidence interval from the 2SLS coefficient.

cost-benefit calculus. Costs, for example, might be overweighted relative to their true value for several reasons. If a patient arrives at the pharmacy counter to find that her drug basket has shot up in cost relative to her expectations, costs may be highly salient (Bordalo, Gennaioli, and Shleifer 2013, 2020); if costs deviate from previously set reference points, they may be viewed as losses (Kahneman and Tversky 1979). Present bias (Laibson 1997; O'Donoghue and Rabin 1999) could likewise cause patients to overweight present costs over future benefits. Alternatively, patients could be relying on heuristics—like filling the most important drug, dropping the most expensive drug—effectively substituting simpler problems for the more difficult full calculation of marginal costs and benefits (Tversky and Kahneman 1974). Or patients could disengage from the cost-benefit calculus altogether, because of inattention or frictions (Handel and Schwartzstein 2018; Gabaix 2019), choice fatigue (Iyengar and Kamenica 2010; Augenblick and Nicholson 2016), or judging the problem as unsolvable and simply giving up (Ackerman and Thompson 2017). This last set of mechanisms in particular could explain the phenomenon of patients choosing to fill none of their medications in response to price increases. Time-specific factors in December may contribute: while the absence of ZIP-level effects suggest that household budgets around the holidays are less likely to explain our results, the administrative burdens of getting a doctor's visit or responses to questions may be larger than usual at this time. Exploring which of these factors might be at play is a fruitful direction for future work in behavioral science, with potentially large real-world impact.

It is also worth considering a simpler explanation: the seriousness of interrupting drugs is simply not known to patients—and potentially their doctors. Our survey of patients taking medications, while small, is some of the first evidence of its kind on how patients view short drug interruptions. We partnered with the firm Survey Healthcare Global to recruit 200 patients ages 61–70 who reported taking at least one prescription medication in July 2022. Survey respondents were selected to be similar to the middle spenders in our sample: the median respondent was on at least five medications, and 75% indicated they had hypertension or high cholesterol. The survey took less than five minutes to

complete and had a 100% completion rate.[45] As shown in Figure I, patients view short interruptions as innocuous: two-thirds doubt any acute events (hospitalization, deaths) would result from even a month-long interruption. Most cannot imagine any issues with missing their drugs for a week. These beliefs are one explanation for a central driver of our results—that drug interruptions peak sharply in December—and may also explain why the phenomenon is so temporary: patients are willing to hold out for low prices just over the horizon in January, because they view short-term interruptions as innocuous. Our results argue strongly that this view is mistaken.

Ultimately, the decision to ingest a drug lies with the patient. However our results suggest that both physicians and policy makers are missing opportunities to improve the architecture of these decisions. Policy makers should remember that drug cost-sharing policies have major implications for patient health, as well as health care costs. Physicians should remind their patients that for a variety of chronic medications, even short interruptions can be deadly.

HARVARD KENNEDY SCHOOL, HARVARD BUSINESS SCHOOL, AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES
STANFORD UNIVERSITY, UNITED STATES
UNIVERSITY OF CALIFORNIA, BERKELEY SCHOOL OF PUBLIC HEALTH AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES

45. The sample was drawn from a panel of over 600,000 patients and caregivers maintained by the firm; more information is available on SHG's website. The exact wording of the questions was: (1) "Patients often miss doses of their medications (research has found that up to 57% of doses are missed). Imagine a situation where you missed doses of your own medications. How long would it take before your risk of a serious health problem increased?" Possible responses were: "<1 week," "1–2 weeks," "2–3 weeks," "3–4 weeks," ">4 weeks." (2) "Think about the kinds of health problems that could arise from missing your medications. Which of the following could happen?" Possible responses were: "no change to your health," "you feel worse on days you miss the medications," "your chronic conditions get worse, in a way that eventually harms your health," "you need to be hospitalized," "death." For this question, 15 respondents selected "no change to your health," and thus do not contribute to the totals in the figure.

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

The data underlying this article are available in the Harvard Dataverse, https://doi.org/10.7910/DVN/F6X52R (Chandra, Flack, and Obermeyer 2024).

References

Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc, "Mortality Effects and Choice across Private Health Insurance Plans," *Quarterly Journal of Economics*, 136 (2021), 1557–1610. https://doi.org/10.1093/qje/qjab017

Ackerman, Rakefet, and Valerie A. Thompson, "Meta-Reasoning: Monitoring and Control of Thinking and Reasoning," *Trends in Cognitive Sciences*, 21 (2017), 607–617. https://doi.org/10.1016/j.tics.2017.05.004

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?," *Review of Economics and Statistics*, 97 (2015), 725–741. https://doi.org/10.1162/REST_a_00518

Augenblick, Ned, and Scott Nicholson, "Ballot Position, Choice Fatigue, and Voter Behaviour," *Review of Economic Studies*, 83 (2016), 460–480. https://doi.org/10.1093/restud/rdv047

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein, "Behavioral Hazard in Health Insurance," *Quarterly Journal of Economics*, 130 (2015), 1623–1667. https://doi.org/10.1093/qje/qjv029

Bain, Kevin T., Holly M. Holmes, Mark H. Beers, Vittorio Maio, Steven M. Handler, and Stephen G. Pauker, "Discontinuing Medications: A Novel Approach for Revising the Prescribing Stage of the Medication-Use Process," *Journal of the American Geriatrics Society*, 56 (2008), 1946–1952. https://doi.org/10.1111/j.1532-5415.2008.01916.x

Bibbins-Domingo, Kirsten, David C. Grossman, Susan J. Curry, Karina W. Davidson, John W. Epling, Francisco A. R. García, Matthew W. Gillman, Alex R. Kemper, Alex H. Krist, Ann E. Kurth, C. Seth Landefeld, Michael L. LeFevre, Carol M. Mangione, William R. Phillips, Douglas K. Owens, Maureen G. Phipps, and Michael P. Pignone, "Statin Use for the Primary Prevention of Cardiovascular Disease in Adults: US Preventive Services Task Force Recommendation Statement," *Journal of the American Medical Association*, 316 (2016), 1997–2007. https://doi.org/10.1001/jama.2016.15450

Blasi, Francesco, Marco Mantero, and Stefano Aliberti, "Antibiotics as Immunomodulant Agents in COPD," *Current Opinion in Pharmacology*, 12 (2012), 293–299. https://doi.org/10.1016/j.coph.2012.01.006

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Memory, Attention, and Choice," *Quarterly Journal of Economics*, 135 (2020), 1399–1442. https://doi.org/10.1093/qje/qjaa007

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience and Consumer Choice," *Journal of Political Economy*, 121 (2013), 803–843. https://doi.org/10.1086/673885

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*, 132 (2017), 1261–1318. https://doi.org/10.1093/qje/qjx013
Buckles, Kasey S., and Daniel M. Hungerman, "Season of Birth and Later Outcomes: Old Questions, New Answers," *Review of Economics and Statistics*, 95 (2013), 711–724. https://doi.org/10.1162/REST_a_00314
Chandra, Amitabh, Evan Flack, and Ziad Obermeyer , "Replication Data for: 'The Health Costs of Cost Sharing'," 2024, Harvard Dataverse, https://doi.org/10.7910/DVN/F6X52R.
Chandra, Amitabh, and Jonathan Skinner, "Technology Growth and Expenditure Growth in Health Care," *Journal of Economic Literature*, 50 (2012), 645–680. https://doi.org/10.1257/jel.50.3.645
Chandra, Amitabh, and Douglas O. Staiger, "Identifying Sources of Inefficiency in Healthcare," *Quarterly Journal of Economics*, 135 (2020), 785–843. https://doi.org/10.1093/qje/qjz040
Chandra, Amitabh, Jonathan Gruber, and Robin McKnight, "Patient Cost-Sharing and Hospitalization Offsets in the Elderly," *American Economic Review*, 100 (2010), 193–213. https://doi.org/10.1257/aer.100.1.193
Choudhry, Niteesh K., Jerry Avorn, Robert J. Glynn, Elliott M. Antman, Sebastian Schneeweiss, Michele Toscano, Lonny Reisman, Joaquim Fernandes, Claire Spettell, Joy L. Lee, Raisa Levin, Troyen Brennan, and William H. Shrank, "Full Coverage for Preventive Medications after Myocardial Infarction," *New England Journal of Medicine*, 365 (2011), 2088–2097. https://doi.org/10.1056/NEJMsa1107913
CMS, "Costs in the Coverage Gap," 2024, https://www.medicare.gov/drug-coverage-part-d/costs-for-medicare-drug-coverage/costs-in-the-coverage-gap
Currie, Janet, and Hannes Schwandt, "Within-Mother Analysis of Seasonal Patterns in Health at Birth," *Proceedings of the National Academy of Sciences*, 110 (2013), 12265–12270. https://doi.org/10.1073/pnas.1307582110
Czosnowski, Quinn A., Joseph M. Swanson, Bob L. Lobo, Joyce E. Broyles, Paul R. Deaton, and Christopher K. Finch, "Evaluation of Glycemic Control Following Discontinuation of an Intensive Insulin Protocol," *Journal of Hospital Medicine*, 4 (2009), 28–34. https://doi.org/10.1002/jhm.393
Doblhammer, Gabriele, and James W. Vaupel, "Lifespan Depends on Month of Birth," *Proceedings of the National Academy of Sciences of the United States of America*, 98 (2001), 2934–2939. https://doi.org/10.1073/pnas.041431898
Dumas, Emily O., and Gary M. Pollack, "Opioid Tolerance Development: A Pharmacokinetic/Pharmacodynamic Perspective," *AAPS Journal*, 10 (2008), 537–551. https://doi.org/10.1208/s12248-008-9056-1
Einav, Liran, and Amy Finkelstein, "Moral Hazard in Health Insurance: What We Know and How We Know It," *Journal of the European Economic Association*, 16 (2018), 957–982. https://doi.org/10.1093/jeea/jvy017
Einav, Liran, Amy Finkelstein, and Maria Polyakova, "Private Provision of Social Insurance: Drug-Specific Price Elasticities and Cost Sharing in Medicare Part D," *American Economic Journal: Economic Policy*, 10 (2018), 122–153. https://doi.org/10.1257/pol.20160355
Einav, Liran, Amy Finkelstein, and Paul Schrimpf, "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D," *Quarterly Journal of Economics*, 130 (2015), 841–899. https://doi.org/10.1093/qje/qjv005
Finkelstein, Amy, and Kathleen McGarry, "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 96 (2006), 938–958. https://doi.org/10.1257/aer.96.4.938
Gabaix, Xavier, "Behavioral Inattention," in *Handbook of Behavioral Economics: Applications and Foundations*, Stefano, B. and David, eds., 2, (Amsterdam: Elsevier, 2019), 261–343. https://doi.org/10.1016/bs.hesbe.2018.11.001
Goldin, Jacob, Ithai Z. Lurie, and Janet McCubbin, "Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach," *Quarterly Journal of Economics*, 136 (2021), 1–49. https://doi.org/10.1093/qje/qjaa029
Handel, Benjamin, and Joshua Schwartzstein, "Frictions or Mental Gaps: What's behind the Information We (Don't) Use and When Do We Care?," *Journal of*

Economic Perspectives, 32 (2018), 155–178. https://doi.org/10.1257/jep.32.1.155

Heeschen, Christopher, Christian W. Hamm, Ulrich Laufs, Steven Snapinn, Michael Böhm, and Harvey D. White, "Withdrawal of Statins Increases Event Rates in Patients with Acute Coronary Syndromes," *Circulation*, 105 (2002), 1446–1452. https://doi.org/10.1161/01.CIR.0000012530.68333.C8

Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica*, 81 (2013), 1713–1762. https://doi.org/10.3982/ECTA10931

Horowitz, Mark Abie, Sameer Jauhar, Sridhar Natesan, Robin M. Murray, and David Taylor, "A Method for Tapering Antipsychotic Treatment That May Minimize the Risk of Relapse," *Schizophrenia Bulletin*, 47 (2021), 1116–1129. https://doi.org/10.1093/schbul/sbab017

Iyengar, Sheena S., and Emir Kamenica, "Choice Proliferation, Simplicity Seeking, and Asset Allocation," *Journal of Public Economics*, 94 (2010), 530–539. https://doi.org/10.1016/j.jpubeco.2010.03.006

Jarad, N. A., J. A. Wedzicha, P. S. Burge, and P. M. A. Calverley, "An Observational Study of Inhaled Corticosteroid Withdrawal in Stable Chronic Obstructive Pulmonary Disease," *Respiratory Medicine*, 93 (1999), 161–166. https://doi.org/10.1016/s0954-6111(99)90001-x

Kahneman, Daniel, and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47 (1979), 263–292. https://doi.org/10.2307/1914185

Kaplan, Cameron M., and Yuting Zhang, "Anticipatory Behavior in Response to Medicare Part D's Coverage Gap," *Health Economics*, 26 (2017), 338–351. https://doi.org/10.1002/hec.3311

KFF, "An Overview of the Medicare Part D Prescription Drug Benefit," 2023, https://www.kff.org/medicare/fact-sheet/an-overview-of-the-medicare-part-d-prescription-drug-benefit/

Koracevic, Goran, Sladjana Micic, Milovan Stojanovic, Miloje Tomasevic, Tomislav Kostic, Radmila Velickovic Radovanovic, Dragan Lovic, Dragan Djordjevic, Miomir Randjelovic, Maja Koracevic, and Zaklina Ristic, "Beta Blocker Rebound Phenomenon Is Important, but We Do Not Know Its Definition, Incidence or Optimal Prevention Strategies," *Hypertension Research*, 43 (2020), 591–596. https://doi.org/10.1038/s41440-020-0449-6

Kulkarni, Ameya S., Sriram Gubbi, and Nir Barzilai, "Benefits of Metformin in Attenuating the Hallmarks of Aging," *Cell Metabolism*, 32 (2020), 15–30. https://doi.org/10.1016/j.cmet.2020.04.001

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112 (1997), 443–478. https://doi.org/10.1162/003355397555253

Miller, Richard R., Harold G. Olson, Ezra A. Amsterdam, and Dean T. Mason, "Propranolol-Withdrawal Rebound Phenomenon—Exacerbation of Coronary Events after Abrupt Cessation of Antianginal Therapy," *New England Journal of Medicine*, 293 (1975), 416–418. https://doi.org/10.1056/NEJM197508282930902

Miller, Sarah, Norman Johnson, and Laura R. Wherry, "Medicaid and Mortality: New Evidence From Linked Survey and Administrative Data," *Quarterly Journal of Economics*, 136 (2021), 1783–1829. https://doi.org/10.1093/qje/qjab004

Mullainathan, Sendhil, and Ziad Obermeyer, "Diagnosing Physician Error: A Machine Learning Approach to Low-Value Health Care," *Quarterly Journal of Economics*, 137 (2022), 679–727. https://doi.org/10.1093/qje/qjab046

Natarajan, Pradeep, Robin Young, Nathan O. Stitziel, Sandosh Padmanabhan, Usman Baber, Roxana Mehran, Samantha Sartori, Valentin Fuster, Dermot F. Reilly, Adam Butterworth, Daniel J. Rader, Ian Ford, Naveed Sattar, and Sekar Kathiresan, "Polygenic Risk Score Identifies Subgroup with Higher Burden of Atherosclerosis and Greater Relative Benefit from Statin Therapy in the Primary Prevention Setting," *Circulation*, 135 (2017), 2091–2101. https://doi.org/10.1161/CIRCULATIONAHA.116.024436

Neumann, Peter J., Joshua T. Cohen, and Milton C. Weinstein, "Updating Cost-Effectiveness—The Curious Resilience of the $50,000-per-QALY Threshold," *New England Journal of Medicine*, 371 (2014), 796–797. https://doi.org/10.1056/NEJMp1405158
O'Donoghue, Ted, and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 89 (1999), 103–124. https://doi.org/10.1257/aer.89.1.103
Oesterle, Adam, Ulrich Laufs, and James K. Liao, "Pleiotropic Effects of Statins on the Cardiovascular System," *Circulation Research*, 120 (2017), 229–243. https://doi.org/10.1161/CIRCRESAHA.116.308537
Oliver, Thomas R., Philip R. Lee, and Helene L. Lipton, "A Political History of Medicare and Prescription Drug Coverage," *Milbank Quarterly*, 82 (2004), 283–354. https://doi.org/10.1111/j.0887-378X.2004.00311.x
Osterberg, Lars, and Terrence Blaschke, "Adherence to Medication," *New England Journal of Medicine*, 353 (2005), 487–497. https://doi.org/10.1056/NEJMra050100
Pauly, Mark V., "The Economics of Moral Hazard: Comment," *American Economic Review*, 58 (1968), 531–537. https://www.jstor.org/stable/1813785
Pollard, John, "Improving Mortality: A Rule of Thumb and Regulatory Tool," *Journal of Actuarial Practice*, 10 (2002), 219–238.
Psaty, Bruce M., Thomas D. Koepsell, Edward H. Wagner, James P. LoGerfo, and Thomas S. Inui, "The Relative Risk of Incident Coronary Heart Disease Associated with Recently Stopping the Use of β-Blockers," *Journal of the American Medical Association*, 263 (1990), 1653–1657. https://doi.org/10.1001/jama.1990.03440120075040
Raffoul, Melanie, Miranda Moore, Doug Kamerow, and Andrew Bazemore, "A Primary Care Panel Size of 2500 Is Neither Accurate Nor Reasonable," *Journal of the American Board of Family Medicine*, 29 (2016), 496–499. https://doi.org/10.3122/jabfm.2016.04.150317
Ridker, Paul M., Eleanor Danielson, Francisco A. H. Fonseca, Jacques Genest, Antonio M. Gotto, John J. P. Kastelein, Wolfgang Koenig, Peter Libby, Alberto J. Lorenzatti, Jean G. MacFadyen, Børge G. Nordestgaard, James Shepherd, James T. Willerson, and Robert J. Glynn, "Rosuvastatin to Prevent Vascular Events in Men and Women with Elevated C-Reactive Protein," *New England Journal of Medicine*, 359 (2008), 2195–2207. https://doi.org/10.1056/NEJMoa0807646
Rosenson, Robert S., "Low Density Lipoprotein-Cholesterol (LDL-C) Lowering after An Acute Coronary Syndrome," *UpToDate*, 2023, https://pro.uptodatefree.ir/Show/100
Steinman, Michael, and Emily Reeve, "Deprescribing," *UpToDate*, 2023, https://www.uptodate.com/contents/deprescribing
Tversky, Amos, and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185 (1974), 1124–1131. https://doi.org/10.1126/science.185.4157.1124
Walma, Edmond, Colette van Dooren, Ad Prins, Emiel van der Does, and Arno Hoes, "Withdrawal of Long Term Diuretic Medication in Elderly Patients: A Double Blind Randomised Trial," *BMJ*, 315 (1997), 464–468. https://doi.org/10.1136/bmj.315.7106.464
Zhang, Yin, Elizabeth E. Devore, Susanne Strohmaier, Francine Grodstein, and Eva S. Schernhammer, "Birth Month, Birth Season, and Overall and Cardiovascular Disease Mortality in US Women: Prospective Cohort Study," *BMJ*, 367 (2019), l6058. https://doi.org/10.1136/bmj.l6058