Title

The Effects of Non-linear Pricing on Consumer Behavior in the U.S. Pharmaceutical Market

Author

Illich, Matthew Kepler

Publication Date

2024

Peer reviewed|Thesis/dissertation

The Effects of Non-linear Pricing on Consumer Behavior in the U.S. Pharmaceutical Market

By

MATTHEW KEPLER ILLICH
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Marianne Bitler, Chair

Erich Muehlegger

Colin Cameron

Committee in Charge

2024

Contents

Abstract

The ACA introduced a regulation in section 1302(c)(1) that requires all private insurance plans to provide customers with a maximum annual out-of-pocket limit. Using a difference-in-differences estimation with continuous treatment and data from the Medical Expenditure Panel Survey and IQVIA, I estimate that the 2015 and 2016 enforcement of this regulation increased the proportion of branded pharmaceutical claims dispensed with no out-of-pocket cost, while no effect was detected for the more limited enforcement in 2014. From 2012 to 2018, the proportion of branded pharmaceutical prescriptions paid with no out-of-pocket costs increased from 8.5% to 20.2% for privately-insured consumers (including employer-provided, health-exchange, and individual plans). This rises to 27.4% in December 2018 as consumers increasingly reach the limit throughout the calendar year. This increase is larger than would have been expected based on the mean pre-ACA exposure estimates which may be evidence of strategic behavior or other policies amplifying the effect of this regulation. I also examine whether exposure to section 1302(c)(1) affected measures of medication adherence and find that it increased medication possession ratio in 2015.

Due to the annual plan-year deductible and out-of-pocket limit reset, there exists a ex-post inequity based on when an adverse medical event occurs for privately-insured consumers in the United States. I use a regression discontinuity design to examine this behavior. For some consumers, the effects are large: for patients who visit the emergency room (ER) with a myocardial infarction, the average coinsurance rates for pharmaceuticals in the following year decreases by 5.4 p.p (21% relative decrease) if the ER visit is in January relative to those who visit in December. Both pricing transitions out of the deductible phase and past the out-of-pocket limit contribute substantially to this differential. These decreases in spot prices do not appear to increase the number of circulatory-system related prescriptions filled, but cause increases in other prescription spending, particularly branded products, providing evidence that myocardial infarction patients rationally prioritize their healthcare spending

when faced with a health crisis. Based on these analyses, the price elasticity of demand for pharmaceutical products for myocardial infarction patients is estimated to be -0.23.

I also examine how patients respond to reaching their out-of-pocket limit. Using a fixed effects regression, I find that once patients reach their out-of-pocket limit, their branded drug consumption increases by 12% and generic consumption grows by 7% in the last four weeks of the calendar year relative to their behavior in years when they do not reach their out-of-pocket limit. Patients residing in the lowest income quartile geographic areas benefit the most, with their December drug utilization increasing four times as much as those in the highest income quartile areas.

Acknowledgments

I am grateful to my graduate advisors, Marianne Bitler and Erich Muehlegger for their helpful feedback throughout the development of this work and their continued guidance. I also thank IQVIA and the people in the Market Access Strategy Consulting group who develop and maintain the LAAD data and shared their expertise and understanding of the data with me. In particular, I thank Rory Martin at IQVIA for his advice throughout the project. Finally, I thank my wife and family for their love and support throughout my PhD program.

CHAPTER 1

Introduction

This dissertation contains three papers exploring the effects of non-linear pricing for privately insured-patients in the United States healthcare market. These nonlinear effects take two primary forms: Deductibles, where patients are required to pay a set dollar amount before the insurer pays for covered medical and pharmaceutical claims, and out-of-pocket limits, where patients are no longer required to pay out-of-pocket for covered care after they have paid a set dollar amount in one plan-year.

The first paper draws a causal relationship between provisions in the Affordable Care Act and a substantial increase in patients reaching their out-of-pocket limit and paying zero dollars for pharmaceutical transactions, especially near the end of the calendar year. I further show that the patients who use pharmaceutical products that were most affected by this legislation increased their adherence (as measured by medication possession ratio) to those products, especially near the end of the calendar year.

The second paper uses the discontinuity in the non-linear pricing that results from out-of-pocket limits and deductibles as a natural experiment to measure the price elasticity of demand for patients who have been admitted to the emergency room and diagnosed with a myocardial infarction (also known as a heart attack). I measure the extent to which having a myocardial infarction at the beginning of the calendar year leads to lower future pharmaceutical prices relative to having a myocardial infarction at the end of the calendar year. I then measure an increase in pharmaceutical purchases, especially branded and non-circulatory medications, for patients who had their myocardial infarction at the beginning of the year relative to those who had it at the end of the calendar year.

The third paper, coauthored by Rory Martin, examines how patients who reach the out-of-pocket limit respond to that change in prices. The focus on this paper is whether patients increase their purchases in the last four weeks of December when the out-of-pocket price is zero and then decrease their purchases in January when the out-of-pocket price resets in the following plan-year. We find that both effects are present, but the increases in December purchasing significantly outweigh the decreases in January leading to higher utilization overall due to the out-of-pocket limit.

CHAPTER 2

Free Drugs in December: The Effects of the Affordable Care Act's Maximum Out-of-Pocket Limit Regulation

2.1. Introduction

Dismay regarding high and increasing healthcare costs in the United States has been a mainstay in the national conversation for the past few decades. These costs are transmitted to citizens through many channels: Government-funded health insurance programs such as Medicare and Medicaid require larger budgets, while uninsured consumers limit care and seek charitable resources or else face insurmountably high medical bills and medication prices (Garthwaite et al. 2018). This paper investigates privately insured consumers who make up about two-thirds of the U.S. population as of 2020 (Keisler-Starkey et al. 2021) and pay for high healthcare costs primarily through monthly premium payments and out-of-pocket (OOP) costs such as deductibles, copayments, and coinsurance. For these patients, the proportion of prescription medications paid for without OOP cost has increased substantially in recent years.

In 2010, The Patient Protection and Affordable Care Act (ACA) was passed with the intention to reduce the number of uninsured Americans, improve healthcare, and lower healthcare costs. Most major provisions in the ACA came into effect in 2014. While many of these provisions have been extensively studied (Gruber and Sommers 2019), one provision that, to my knowledge, has not been studied in the academic literature was implemented in section 1302(c)(1) that required all ACA-compliant private health insurance policies sold in the United States to provide an annual maximum out-of-pocket payment limit for an individual's or family's medical expenses. This annual out-of-pocket limit was initially set

at a maximum of $6,350 for individuals or $12,700 for families of two or more people and increases each year in tandem with premiums as specified in section 1302(c)(4). Figure 2.11 contains the language of the law.

The direct effect of a binding annual OOP limit is to shift the healthcare costs for privately-insured individuals from out-of-pocket expenses to premium payments. This transfer of costs has at least four primary economic effects on consumers: Firstly, it may induce moral hazard and increase healthcare utilization as OOP costs decrease. Secondly, it may reduce the financial risk of health shocks. Thirdly, in combination with regulations on setting premiums, it may transfer costs from chronically sick patients to patients with low health costs. Finally, in combination with regulations on minimum coverage requirements, it institutes a "quality floor" on private health insurance policies. In addition to transferring out-of-pocket costs to premiums, the limit could also indirectly affect bargaining between insurers and providers. The OOP limit reduces the control insurers have over their members' spending which may theoretically reduce their bargaining power and increase the prices they pay for healthcare goods and services.

The requirement of a OOP limit changes patient incentives and may induce moral hazard in consumers who have reached their annual OOP limit and in those who expect they may reach the limit in the future. If these consumers increase their utilization relative to the counterfactual, then this provision could lead to an increase in healthcare consumed and therefore overall costs. The limit may also change the incentives of healthcare providers and insurers, amplifying or attenuating the increase in costs respectively. There has been a large body of work has focused on quantifying how patient costs affect medical spending. However, as (Aron-Dine et al. 2015) explore, patient costs are generally non-linear due to the design of health insurance contracts. The two most common ways that non-linearity occurs is through deductibles and OOP limits. In their paper, they find that the initial utilization of healthcare is affected by the future price as well as the spot price. They make the prescient

remark that there is no one true "demand elasticity" since it depends on where consumers are located in their annual cost curve. The current work contributes to this line of research by showing that the ACA substantially increased the non-linearity of insurance contracts in private insurance markets though the OOP limit.

While transferring costs from uncertain OOP costs to predetermined premiums may induce moral hazard, it reduces financial risk to patients in the event of negative health shocks. This reduction in the financial risk of health shocks may affect the prevalence of bankruptcies. The literature on medical bankruptcies in the United States is mixed in terms of the magnitude with a recent estimate finding that hospitalizations only cause modest increases in bankruptcies (Dobkin et al. 2018) while previous studies show that medical problems may have contributed to a large amounts of bankruptcies (Himmelstein et al. 2009) (Hackney et al. 2016). Since some healthcare costs are so high that uncapped percentage-based coinsurances can be very onerous to consumers, a maximum OOP limit may be integral in the policy goal of reducing medical financial stress and medically related bankruptcies.

The effect of the OOP limit on equity stems in part from other sections of the ACA limiting insurance plans from charging differential monthly premiums with the exception of five factors which can affect a plan's monthly premium: geography, age, tobacco use, plan category, and whether the plan covers dependents. If some individuals are very unlikely to reach the OOP limit, they are required to pay for that coverage and do not receive any benefits in terms of reduced premiums for being unlikely to reach that limit. Alternatively, if other individuals have pre-existing conditions and are more likely to reach the OOP limit, their expected OOP costs decrease in the end of the plan year, and they do not pay for that increased coverage with increased premiums relative to individuals without pre-existing conditions. The combination of regulations requiring both non-discriminatory premiums and minimum required benefits transfer wealth from those who are expected to be healthy to those who are not.

Finally, this provision is necessary to create a regulated "quality floor" of ACA-compliant private health insurance in that the law specifies which services insurers must cover, and with section 1302(c)(1), a cost-sharing maximum. There is a large academic literature that supports the idea that many healthcare consumers may be unable to make consistent health insurance decisions without assistance (Bhargava et al. 2017) (Abaluck and Gruber 2011). So a health insurance "quality floor" with a clearly delineated limit on cost exposure may help consumers if they do not have the ability or information to accurately compare insurance plans with more complex cost-sharing schemes.

As far as I am aware, no economics research papers have specifically focused on the effect of the maximum OOP limits included in the ACA in 2014 through 2016. This dearth in research may have two explanations: Firstly, OOP cost data covering privately insured patients can be difficult to obtain since the only potential data sources are insurer administrative data, relatively small surveys, data from individual private firms, or pharmacy claims data. Secondly, there may be a preconception that this part of the ACA was not very impactful relative to the other major provisions. However, the inclusion of $2,000 OOP limits for Medicare drug costs in the Inflation Reduction Act of 2022 may increase the importance of understanding this policy.

This paper investigates the direct impact of the maximum OOP limit instituted in section 1302(c)(1). I look to answer two related questions: Did this change in the law increase the proportion of privately-insured patients who pay zero out-of-pocket cost for branded prescription medications? If so, how much of the impact was related to certain changes in the enforcement of the law between 2013 and 2016? I also evaluate impacts on medication adherence (medication possession ratios) for branded and generic pharmaceutical products. I approach these questions using a difference-in-differences model with continuous treatment across branded pharmaceutical products based on the pre-ACA product-level OOP cost distribution. This leverages the variation in products that are used by patients with high

expected annual OOP costs compared to products used by patients with low expected annual OOP costs.

In both the Medical Expenditure Panel Survey (MEPS) and IQVIA data sources, I find descriptive evidence of a large increase in zero-dollar OOP costs for patients using branded pharmaceutical products between 2013 and 2016 coinciding with the enforcement of section 1302(c)(1). I then estimate the pre-ACA exposure to the maximum OOP limit for branded products using MEPS out-of-pocket total and drug spending data. I then leverage differentials between products in exposure to this regulation to estimate the direct impact of the law and to provided evidence that this regulation is driving the increase in zero-dollar OOP cost prescriptions. However, I find that the pre-2014 exposure estimates are substantially lower than the descriptive increases in zero-dollar OOP prescriptions.

2.2. Institutional Background

The expected effect of this provision depends on three factors: Firstly, it depends on the extent to which this section in the ACA influences insurance contract design. If many insurers already had an OOP limit less than the mandated maximum, then there would be no effect on those plans. Alternatively, the ACA allowed insurers to maintain "grandfathered" health plans that are not covered by this restriction. According to the 2019 Employer Health Benefits survey, about 13% of workers were enrolled in grandfathered plans (Kaiser Family Foundation 2019). For these workers, the mandated maximum would have no effect. It is also possible that insurance companies were unable or unwilling to comply with this regulation. Additionally, insurers may take advantage of an exception in the law that only requires cumulation of "in-network" OOP costs towards the mandated maximum OOP limit. Secondly, the effect depends on the extent to which consumers pay more than the OOP limit. If very few consumers ever reach $6,350 in OOP costs in the year or do so very late in the year, then the impact of the law will be muted even if health insurance contracts are being affected by the law. Finally, the effect may be amplified or attenuated by the agents' strategic

responses to the regulation. For example, pharmaceutical manufacturers could raise prices or increase coupon prevalence to drive patients to the limit faster, thereby reducing demand elasticity for their products. Likewise, patients could increase spending early in the year to take advantage of zero-OOP cost care later in the year. Insurers may also modify out-of-pocket expenditure schedules (copayment and coinsurance amounts) or change bargaining behavior with healthcare providers and pharmaceutical manufacturers.

To understand whether section 1302(c)(1) is likely to be impactful, it helps to examine the institutional background before the law went into effect in 2014. Prior to the ACA, private insurance plans were not required by law to provide policyholders with an OOP limit. Despite this, 88% of employer-sponsored plans had an OOP limit and 43% had OOP limits of less than $3,000 in 2013 (Kaiser Family Foundation 2013). These high levels of OOP limits may support the idea that the maximum OOP limit section may not have been binding to many employer-sponsored plans. However, before 2014 some costs such as physician visit co-pays were not included as contributing to the OOP limit. Furthermore, plans that did offer OOP limits may have had separate OOP limits for medical care and pharmaceutical care. This made it more difficult to reach either limit so many patients did not often reach the limit, even if their plan had one. Finally, the individual private insurance market, which insures about 18 million Americans, may not have had the same proportion of plans with OOP limits as the employer-sponsored insurance market that is surveyed by (Kaiser Family Foundation 2013), so it is possible that the law may have had a larger impact in individual markets than in employer group markets.

In addition to pre-ACA OOP limits already in place, it is possible that very few privately-insured individuals have enough utilization and spending to reach OOP costs near the mandated limits. This would prevent the regulation from binding on many individuals. Indeed, this appeared to be the case in 2013 when less than 1% of private group insurance beneficiaries had OOP costs higher than $6,350 (Glied and Zhu 2020). It is possibly for this reason

that this regulation has not been studied by other researchers. However, many participants in the healthcare industry deal almost exclusively with the part of the patient distribution with the highest OOP costs. For instance, a dialysis or oncology clinic is not likely to deal with someone in the bottom half of the OOP distribution. Many healthcare providers, pharmaceutical manufacturers, and pharmacies all interact with disproportionately high-cost patient populations. So from their perspective, the proportion of their patients paying $0 after reaching the OOP limit may be expected to substantially exceed 1%. To put it concretely, a larger share of the potential customers in high-need healthcare markets may be on the margin of reaching a mandated OOP limit than the general population.

Before the ACA went into effect, there had been a trend of rising deductibles and spending on coinsurance. For instance, (Kaiser Family Foundation 2018) finds that the average spending on deductibles increased by 80% from 2004 to 2009 before the ACA was passed, and by an additional 81% from 2009 to 2013, the year before section 1302(c)(1) went into effect. Over the same time periods, coinsurance spending outpaced both the growth in costs paid by private insurance as well as workers wages. Overall, average total OOP cost sharing increased from $422 per year in 2004 to $742 in 2013. These trends are important to this analysis because as OOP cost sharing increases, policyholders are more likely to reach the OOP limit given the same medical or pharmaceutical utilization level. Subsequently, if these rising cost trends were to continue in combination with the ACA maximum OOP limit statute, then more patients may reach the OOP limit each year than would be estimated based on exposure to OOP costs in the pre-implementation period.

The government sponsored insurance programs were unaffected by this change: Medicare Part D and traditional Medicare continue to have no OOP limit, while Medicare Advantage has been regulated to include OOP limits of no more than $6,700 since 2011. So these large segments of the U.S. healthcare market are not evaluated in this study. One might be tempted to consider them as potential control groups to compare against the private

insurance market, but the patient populations are substantially different and the insurance contracts lead to different OOP cost schedules for policyholders. For instance, in Medicaid, co-payments are very small relative to private insurance co-payments and patients have lower incomes than those with private insurance ($8 was the maximum allowable Medicaid co-payment for a prescription drug in 2015 while copayments in private insurance plans range from an average of $11 to $93 (Kaiser Family Foundation 2015)) . Meanwhile, Medicare Part D has a complex OOP cost structure that leads to OOP costs that are somewhat comparable in magnitude to private insurance, but with distinct phases of coverage that are unlike private insurance (particularly the Part D coverage gap) and with an older population than the privately-insured population.

The rollout of this regulation had three distinct phases: Firstly, due to technical considerations of merging medical and pharmacy data systems, health insurance providers were allowed a regulatory safe harbor in 2014 wherein, if they had multiple vendors for their medical and pharmaceutical plans, then the statutory OOP limits were separate for the 2014 calendar year (HHS 2013). There was also an additional loophole in 2014 wherein if a prescription drug plan did not have an OOP limit and the insurer had separate vendors for their medical and pharmaceutical plans, then the maximum OOP limit regulation only applied to medical costs. Because of this, patients would be expected to be less likely to reach the OOP limit in 2014 than in subsequent years. By 2015, the regulatory safe harbor policy concluded and all private health insurance plans were required to abide by the statutory annual OOP limits for a combined prescription drug and medical spending category. In 2016, the Department of Health and Human Services tightened the rules for family plans by requiring that both the individual OOP limits and family OOP limits bind to individual family members (HHS 2015). For example, if in 2016 an individual in a family plan spends more than the individual OOP limit, then any expenses for that family member above the individual OOP limit should be paid for by the health insurer. So after January of 2016, an

individual in a family is protected by both the individual OOP maximum for themself and by their family OOP maximum for their family.

By 2019, the ACA individual mandate no longer applied at the federal level. The mandate may have been integral in enforcing the maximum OOP limit because consumers no longer have as much incentive to purchase ACA-compliant healthcare coverage. Because of this, the data in this study do not go past December 2018. Future work may be valuable in understanding how the removal of the individual mandate affects this and other aspects of the ACA insurance requirements.

2.3. Data

2.3.1. Data and Inclusion Criteria. The IQVIA Longitudinal Access and Adjudication Data (LAAD) anonymized pharmacy claims dataset contains claim-level information on a substantial proportion of pharmacy transactions in the United States between the years of 2012 and 2018. There are several key features of this data that make it well-suited for answering my research questions. Firstly, all claims can be attributed to an anonymized patient identifier that allows me to track patients' claims over the course of several years. Claims can also be attributed to a primary payer (often an insurer or pharmacy benefit manager) and a secondary payer if one exists. Common secondary payers are manufacturer provided coupons or supplemental insurance. Secondly, for many claims, the LAAD dataset contains information detailing the cost of each transaction and the portions paid by each paying party (usually some combination of insurers, patients, and/or discount programs). This allows me to identify the payment channel each for each claim so that I can separate Medicaid, Medicare, cash, and private insurance transactions. Thirdly, the data contains the OOP costs both before and after manufacturer-provided coupons. For this study, I focus on the pre-coupon OOP costs because these are the costs that are accrued to the OOP limit in the study period. The data also contains information relating to each individual

prescription. Some of the information I use in this study is related to the product dispensed, the quantity dispensed, and the number of days over which the prescriber intended the prescription to be used. The UC Davis Institutional Review Board determined that this research using anonymous patient data was not human subject research (IRB protocol number: [1920059-1]).

Table 2.2 shows descriptive data from the IQVIA sample of branded prescriptions used to estimate the proportion of prescriptions dispensed at zero dollars. I focus my analysis on a group of branded medications that each had over $10 million dollars in gross revenue in 2012. This criteria excludes products that were launched after 2012 or did not see substantial uptake during 2012. I also limit the products to exclude contraceptives and smoking deterrent products as their OOP costs were directly affected by other regulations of the ACA as well as other over-the-counter products and those not commonly covered by private insurance. This excluded group of products consists of baby care products, infant formula, insect repellents, nutrients and supplements, diagnostic aids, and various non-drug products. This results in a panel of 734 unique branded products included in the analysis before applying sample constraints from MEPS. Despite these restrictions, the branded products panel contains 98% and 61% of the dollar-denominated sales in 2012 and 2018, respectively. The decrease in 2018 is largely due to new product launches in the intervening years. An additional panel of generic medications without branded competition in 2012 was created for robustness analyses that excluded the same product groups as above as well as products with under 500,000 days of therapy captured in the IQVIA LAAD data in 2012.

The data covers a large proportion of the U.S. branded healthcare market with claims filed by over 50 million consumers filled at 28,192 pharmacies in 2012 and 48,230 in 2018. This increase in pharmacies in sample is a result of increased data capture over the intervening years. All major payment channels are covered in the data. These distributions do change substantially over the course of the study with Medicaid increasing substantially as

a proportion of all prescriptions and private insurance decreasing. This is likely due to a combination of the impacts of the ACA and a change in the composition of brands due to products that lost exclusivity between 2012 and 2018. The demographics of the privately-insured patients throughout the study are relatively consistent and adults under age 65 make up the majority of the sample. Furthermore, some of the included products lost exclusivity during the studied time period. For those products that lose exclusivity, the originator brands are included in the analysis but generic versions of the drug are not included.

I also use the Medical Expenditure Panel Survey (MEPS) to measure medical and pharmaceutical out-of-pocket costs for patients. The MEPS is a nationally representative survey of families and individuals, their medical providers, and employers across the United States. This data source, while limited in sample size compared to the LAAD data (30,000 to 40,000 individuals are interviewed each year), provides representative and comprehensive medical and pharmaceutical OOP cost data for the surveyed households. I use data from 2008 to 2013 to identify the pre-ACA expected OOP cost for the patient populations of each branded drug that was on the market in 2012.

MEPS data is also used to understand how representative prescription drug purchasers are of the general United States privately-insured population. Of the MEPS privately-insured patients under 65 years old, which comprise 55.8% of the survey-weighted MEPS respondents, 62% had at least one prescription of any kind in 2012. 24.5% had at least one prescription in the panel of branded products and 53.7% had at least one prescription in the panel of generic molecules. However, by 2018, only 49.1% filled at least one prescription in the generic molecule panel and 13.5% filled at least one prescription branded product panel due to losses of exclusivity and the introduction of new competing products in the intervening years. See Table 2.3 for more details. Overall, privately-insured individuals who use brands are a minority, but a large and important one from the perspective of pharmaceutical manufacturers, PBMs, health insurers, and pharmacies.

2.3.2. Trends in Zero-Dollar OOP Prescriptions. I first examine the branded market-wide trend in the proportion of prescriptions dispensed to privately-insured patients at zero OOP cost. Figures 2.1 and 2.2 compare the MEPS and IQVIA data for branded and generic products respectively. The 95% confidence intervals [1] are shown for the MEPS data while the IQVIA data has over 75 million prescriptions in each year so the confidence intervals are not visible. The two datasets both show that the proportion of zero-dollar OOP prescriptions increased from 2010 to 2018 and the IQVIA proportion lies near the 95% confidence interval for all years except for 2016 when the MEPS proportion of zero OOP branded claims is below the IQVIA measurement. However, the IQVIA trend appears to be smooth throughout the years and covers administrative data from a substantial proportion of the population of prescriptions. The concordance between the representatively sampled MEPS data and the IQVIA LAAD data gathered from pharmacies and third-party administrators supports the idea that the LAAD data is representative of the US market at least in regards to the proportion of zero-dollar branded OOP claims. For generic products, the LAAD data shows a consistently higher proportion of generic prescriptions paid for at zero OOP cost than does the MEPS data however the trend over time is consistent in the two data sources.

While Figure 2.1 shows that zero-dollar branded prescriptions increased dramatically between 2010 and 2018, it did not prove that these were due to patients reaching their OOP limits. There are many reasons why a drug might be free to the patient: workers' compensation often pays for free-to-the-patient medication for injuries sustained on the job. Another potential reason for zero-dollar copays would be if insurers set zero-dollar copays on especially cost-effective treatments to incentivize patients to utilize these cost-effective treatments and avoid more costly alternative treatments or health outcomes. These incentives would be most likely found among highly cost-effective generic medications which are excluded from the branded analysis but may be common in the generic analysis. Finally,

[1]The MEPS standard errors are calculated accounting for the survey weightings while the IQVIA standard errors are unweighted.

insurers could offer plans that do not require any patient cost-sharing for branded or generic prescriptions (a full insurance plan).

One aspect that distinguishes OOP-limit zero-dollar prescriptions from these other potential zero-dollar prescriptions is that OOP-limit zero-dollar prescriptions are much more likely to come at the end of the insurance plan year as patients build up OOP spending throughout the course of the plan year. This prompts the question: When do the insurance plan years reset? Kaiser Family Foundation's Health Benefits Annual Survey reports that 69% of covered workers by firms offering health benefits have their plan year begin in January (Kaiser Family Foundation 2017a). The next most common months are July with 6% and December and October with 4% each. However, my analysis of the IQVIA LAAD data from 2014 to 2018 at the plan level showed that out of all plans for which I could positively identify a month when the proportion of zero-dollar claims decreased, 87% to 94% (depending on the year) of the volume occurred in January. This is indicative that the vast majority of plans reset their OOP limit in January. This discrepancy between Kaiser's survey and the LAAD data could be related to three factors. Firstly, Kaiser's survey does not include non-employer group and individual insurance plans whereas the IQVIA LAAD and MEPS data do. The non-group plans may be more likely to reset in January than employer plans because that is when policies obtained during open enrollment begin the plan year for the health exchanges. Secondly, it is possible that plan years reset at a different time than when their associated OOP limits resets. (HHS 2015) implies that this is allowed. Thirdly, it is possible that IQVIA cannot differentiate insurance plans as precisely as can Kaiser in their survey. If larger plans with January plan year start dates were mixed with smaller plans with other plan year start dates, my analysis may classify the whole group of plans as January start date. The individual markets and differential reset timing factors would lead to Kaiser underestimating the proportion of OOP limits resetting in January, while the cross plan aggregation would lead to IQVIA LAAD analysis overestimating the proportion of plan

years starting in January. So the true proportion of private insurance plan years starting in January is very likely somewhere between 70% and 90%. This is a large enough proportion that it can expected that if OOP limits are driving the zero-dollar prescriptions, then the zero-dollar prescriptions will be most prevalent in December and least prevalent in January.

In Figure 2.3 I measure the zero-dollar OOP proportion by month and year and find that zero-dollar prescriptions are more likely in December than January in all years. Furthermore the December-January differential grows over the years. This may be due to a combination of the section 1302(c)(1) maximum OOP limit increasingly binding, or an increase in the non-group market due to the health exchange plans (therefore a higher proportion of plans resetting in January), or spillover effects from the individual market driving employer plans to reset in January as well. Regardless of the driving forces, the fact remains that between 2012 and 2018, the proportion of zero-dollar OOP prescriptions have become more cyclical throughout the course of the year. In January 2018, 12.4% of prescriptions were zero-dollar OOP and in December 2018, 27.4% of branded prescriptions were zero-dollar OOP. This compares to 2012 with 7.1% and 11.4% respectively. Over the past several years, the US privately-insured pharmaceutical market has developed an annual cycle of drastically varying patient OOP costs decreasing throughout the calendar year then jumping up as the plan resets into the next calendar year. Finally, to the extent that not all plans reset in January, the price pass-through cyclicality for an individual patient may be even greater than these averages suggest.

2.3.3. Analysis of Pre-ACA OOP Spending. Before the enactment of the law, policy makers may have observed that the proposed maximum OOP limit of $6,350 would bind on relatively few individuals. Using data the 2011 to 2018 MEPS, Figure 2.4 shows the distribution of total OOP healthcare costs for a nationally representative sample of privately-insured consumers who filled at least one branded prescription throughout the course of the

year. As expected, the distribution is substantially right-skewed, and only 1.96% of privately-insured branded pharmaceutical consumers in 2013 spent over the inflation-adjusted OOP limit threshold that would be implemented in 2014. Only consumers who spend over this amount (plus the statutory annual increases to compensate for inflation) in subsequent years will be directly affected by the maximum OOP limit provision. However, this proportion spending over the OOP limit has been relatively consistent both before and after the ACA implementation indicating that either (1) the OOP limit is not impactful potentially because of out-of-network expenses or (2) the MEPS self-pay data that I gather is somehow not observing patients reaching the limit. A similar trend is present in Figure 2.5 which shows family OOP spending.

Table 2.1 aggregates the expected exposure across products based on 2008 to 2013 MEPS OOP cost data and calculates the prescription-weighted proportion of patients who purchase these branded products that would be affected by the section 1302(c)(1) regulation. 100% represents the inflation-adjusted $6,350 limit in 2014 and the 80% and two-thirds rows indicate the average exposure had the limit been set at a lower fraction of the actual limit. For all three limit ratios, the 2016 guidance exposes about 50% more patients to the limit than does the 2015 guidance. This provides a prediction that the pre-ACA to 2016 impact should be about 50% larger than the pre-ACA to 2015 impact if this provision is driving all of the increases in zero-dollar OOP prescriptions. Indeed, this prediction is consistent with the results in Figure 2.3. The exposure to the 2014 change for drug-only expenditures was calculated to be 1.0%. This is a smaller proportional impact than was observed for 2014 in Figure 2.3. It is possible this discrepancy is due to some insurers having joint medical and pharmacy claims systems prior to 2014 and therefore preemptively complying with the 2015 regulation in 2014.

Overall, this descriptive data shows that zero-dollar OOP prescriptions are becoming increasingly common in private insurance and provides circumstantial evidence based on the

timing and magnitudes of the increases that section 1302(c)(1) of the ACA may be related to the trend, but it does not provide direct causal evidence of such. The MEPS OOP cost distribution data between 2011 and 2018 does not provide any evidence of a decrease in weight of OOP cost distribution beyond the legislated limits. The next section will approach the estimation more rigorously using difference-in-differences across branded products with continuous treatment related to the pre-ACA exposure to section 1302(c)(1) for each brand. Through this analysis, I develop a causal link between high pre-ACA OOP costs and the increase in zero-dollar OOP prescriptions.

2.4. Empirical Strategy

The ACA affected the whole U.S. healthcare landscape, so obtaining a causal estimate of the impact of the ACA maximum OOP limit on the proportion of zero-dollar prescriptions is plagued by the standard applied economics problem of an insufficient control group. An ideal group would be a representative cohort of insurance plans that for some reason were not held to this regulation, but were held to all other aspects of the ACA. Alternatively, I could look for products for which no patients would ever have high enough costs to reach the legislated maximum OOP limit. While perfect control plans and products are not available, there is a wide range in the distribution of total OOP cost across products. The patient populations using some products (those with high costs associated with the disease or many comorbidities) are much more likely to be affected by the regulation, and those populations using other products with less healthcare costs associated are less likely to be affected by the regulation. So my primary identification strategy uses a difference-in-differences with continuous treatment framework leveraging the differences in OOP cost distributions across pharmaceutical products (Angrist and Imbens 1995) (Callaway et al. 2021). These OOP cost distributions interact with the regulation in each phase of the legislation. I use cross-product variation in the distribution of total (medical and pharmaceutical) OOP costs and pharmaceutical only OOP costs to estimate the exposure to the policy regimes in 2014,

2015, and after 2015. The 2014 exposure depends on total drug costs (excluding non-drug health costs) since insurers were given a reprieve in this year to integrate their medical and pharmacy systems. In 2015, the exposure relative to the pre-ACA time period depends on the total OOP cost distribution since all plans were required to merge systems into a single OOP limit by this year. However, individuals in family plans had to reach the family maximum limit in order to reach the OOP limit. In 2016 and thereafter, individuals in family plans could reach the OOP limit either by reaching the individual limit or the family limit. I use the MEPS data to estimate total OOP costs as well as pharmacy OOP costs and the IQVIA LAAD data as an alternative estimate of pharmacy OOP costs.

I estimate the following equation:

$$\begin{aligned} Y_{jy} =& \beta_0[Expos2014_j] \times \mathbb{1}\{y = 2014\} + \beta_1[Expos2015_j] \times \mathbb{1}\{y = 2015\} + \\ & \beta_2[Expos2016_j] \times \mathbb{1}\{y > 2015\} + \delta_j + \delta_y + \epsilon_{jy} \end{aligned} \tag{1}$$

where Y_{jy} is a measure of zero-dollar OOP proportion or medication adherence for product j in year y. In my main annual specifications, the two dependent variables I focus on are the proportion of prescriptions paid at zero OOP cost in December and throughout the full calendar year. $Expos2014_j$ is the script-weighted proportion of patients who fill drug j that have annual drug OOP costs over the 2014 limit in 2013 IQVIA data. This is the estimate of the proportion of prescriptions for product j that would be filled after reaching the OOP limit based on the 2014 rules. $Expos2015_j$ is the script-weighted proportion of patients who fill product j that have annual total (medical + drug) OOP costs over the limit in 2008 to 2013 MEPS data. This is the estimate of the proportion of prescriptions for product j that would have been filled after reaching the OOP limit based on the 2015 rules. $Expos2016_j$ modifies the 2015 exposure by applying the OOP limit for individuals within families to both the family and the individual OOP limits rather than just the family limit as was done in 2015. $\delta_j + \delta_y$ are product and year fixed effects.

I also estimate a regression at the calendar month level:

$$
\begin{aligned}
Y_{jt} = &\sum_{m=2}^{12} \beta_{0m} \mathbb{1}\{M(t) = m\} \times [Expos2014_j] \times \mathbb{1}\{Y(t) = 2014\} + \\
&\sum_{m=2}^{12} \beta_{1m} \mathbb{1}\{M(t) = m\} \times [Expos2015_j] \times \mathbb{1}\{Y(t) = 2015\} + \\
&\sum_{m=2}^{12} \beta_{2m} \mathbb{1}\{M(t) = m\} \times [Expos2016_j] \times \mathbb{1}\{Y(t) > 2015\} + \delta_j + \delta_t + \epsilon_{jt} \qquad (2)
\end{aligned}
$$

In the monthly specification above, the dependent variable, Y_{jt}, is the proportion of zero-dollar prescriptions for product j in month t where $M(t)$ is the calendar month and $Y(t)$ is the calendar year of time t. The month of January is left out since it is the least likely month to be affected by the OOP limit. The exposure variables are defined in the same way as in Equation 1.

This method is valid for identification of the causal effect of the constructed estimates of exposure to section 1302(c)(1) on the proportion of zero OOP costs prescriptions in December if the following identifying assumptions hold: (1) no products are treated in the pre-period before 2014 and some products are unexposed in the post-period after 2015, (2) there is no anticipation and outcomes are observed, (3) strong parallel trends: for all exposures to the regulation, the average change in the proportion of zero OOP cost prescriptions for all products would be the same if they had been assigned to that exposure (Callaway et al. 2021). Assumption 1 is not in danger since the regulation was not in effect before 2014. While many plans had OOP limits in the pre-period, this research is seeking to understand the effects of the regulation on the proportion of patients reaching the OOP limit and obtaining free prescriptions. Patients in these plans could be viewed as "always takers" in the potential outcomes framework. Assumption 2 is similarly satisfied at the product level since manufacturers have little control over pre-coupon OOP costs. Patients could modulate their

behavior in 2013 based on the expectation that they would reach a more attainable OOP limit in 2014 or 2015, but this would require significant sophistication in understanding the ACA which included many complicated changes and multi-year planning in 2013 to look forward to 2015 since the 2014 exposure estimate was small due to the reduced enforcement in that year. Insurers on the other hand could have transitioned their insurance contracts towards the ACA OOP limits in the pre-period to get ahead of compliance. However, the effect of this would bias the difference-in-differences methodology to underestimate the impact of the regulation due to this behavior. So while this may be a consideration in interpreting the magnitude of my findings, it should not lead to a false-positive result. Assumption 3 is the fundamentally untestable assumption that plagues difference-in-differences methodologies since we cannot observe how untreated units would have responded to treatment. However, I run parallel trends analysis in the pre-treatment 2012 to 2013 period and cannot reject that the pre-trends are parallel conditional on exposure to the regulation (See Appendix).

2.5. Main Results

2.5.1. Zero OOP Cost Proportion. Table 2.5 presents the results of the difference-in-differences estimate of the effect of pre-ACA patient OOP cost distributions across products on the proportion of branded prescriptions with zero OOP costs in December and year-round. The coefficients of interest relating to the interaction terms between the proportion of patients with OOP spending over 100% of the OOP limit imposed by the ACA and the years when the regulation was enacted can be interpreted as the effect of an additional percentage point of a product's patient population being exposed to OOP costs over the mandated limit on the proportion of zero-dollar OOP claims for that product. The significant and positive results for the 2015 and 2016 changes indicate that the proportion of patients whose OOP costs exceeded the limits before the ACA led patients to reach the OOP limit once those policies were in effect. The non-significant result for the 2014 change indicates that the pre-ACA drug OOP costs were not predictive of the increase in zero-dollar prescriptions in

December. This non-significance may be evidence that the drug-only OOP limit was not very impactful; either because few insurance plans had drug-only OOP limits and the law did not apply to those plans that did not have a drug-only OOP limit or because the limit was too high to bind for many patients. The results for generic products in Table 2.6 are directionally similar to the branded prescriptions.

As expected, for both branded products and generic molecules, the measured coefficients are smaller for the full year than they are for December since many consumers in plans that reset in January do not reach the limit until the end of the calendar year. Figure 2.8 shows the results for the monthly model leaving out January. The coefficients of this regression are statistically significant at the end of the year for the 2015 and 2016 policy changes, while the 2014 coefficients hover around zero effect. This is aligned with the results from Tables 2.5 and 2.6 and controls for biases that would be expected to affect the beginning and the end of the calendar year equally. The products comprising the generic molecules have much lower average prices than the branded products so any OOP costs directly due to the purchase of the product are small for generics. Even so, the effects are still statistically significant for generic molecules. This data provides evidence that the policy changes in 2015 and 2016 requiring insurance plans to have merged OOP limits between medical and drug coverages appeared to bind differentially across both brand and generic products based on their exposure to these changes, especially in the end of the calendar year when more patients will be expected to have reached the OOP limits.

In an idealized world where OOP costs are static across years, all plans reset on January 1st, and all regressors are precisely measured, it would be expected that for every patient exposed to OOP costs over the limit, an additional patient would pay zero dollars in late December. In that ideal case, the interaction terms would be equal to 1. However, these assumptions are not realistic: it is possible the OOP cost distributions changed between 2008 and 2018 for many products, I estimate that 10% to 30% of private insurance plan

years do not start in January, and due to sample limitations in MEPS, the exposure is imprecisely measured for small volume products. Still, the evidence that Section 1302(c)(1) drove a substantial proportion of the increase in zero-dollar OOP prescriptions is strong. The overall increase in zero-dollar OOP branded prescriptions purchased in December between 2013 and 2015 was 8.7% of which 3.6% can be explained by differences in exposure between branded products. Likewise, the overall increase in zero-dollar OOP branded prescriptions purchased in December between 2013 and 2016 was 14.0% of which 5.0% can be explained by differences in exposure between products.

2.5.2. Medication Adherence. From the previous section, exposure to OOP limits was correlated to an increase in the proportion of zero OOP prescriptions at key implementation points of the ACA in 2015 and 2016. A literature of healthcare economics research on moral hazard and price elasticity of healthcare has shown through controlled experiments like the RAND experiment and Oregon Health Insurance Experiment that U.S. consumers respond to prices in their healthcare purchases and purchase more when health insurance is cheap and less when it is expensive. A parallel line of research in health policy has provided evidence that high OOP costs negatively affect drug adherence (Karter et al. 2018) (Curkendall et al. 2008) (Hung et al. 2021) (Kirkman et al. 2015) (Doshi et al. 2009). Poor drug adherence often leads to negative health outcomes since some drugs, especially generic medications, are some of the most cost-effective interventions for many conditions (Shrank et al. 2011) (Neumann et al. 2000).

Does this increase in zero OOP prescriptions translate into increased medication adherence and utilization of these drugs relative to those less exposed drugs? This question could be explored with patient-level longitudinal data in future research, but it is possible to extend the product-level regression from Equation 1 to understand the impacts of Section 1302(c)(1) on adherence. The main empirical strategy of difference-in-differences with continuous treatment effects in the form of exposure to the OOP limit was applied using

measures of adherence as the dependent variable. Three measures are presented here: (1) Calendar Year Medication Possession Ratio (MPR) was calculated by tracking patients who used a product within the first three calendar months of a year from the day they filled their first prescription and for 365 days thereafter. The total days of therapy as indicated by the physician on the prescriptions filled was measured and divided by 365. (2) Fourth quarter days filled was calculated by examining the same patient-product combinations that were on therapy in the first three calendar months of a year and evaluating how many days of therapy they filled in the three months between October and December of the same calendar year. This is the time of the year in which the OOP limits are expected to be most impactful on adherence when the largest proportion of patients are paying for prescriptions at zero OOP cost. (3) Finally, the difference in days of adherence was calculated between the fourth quarter and the following first quarter of the next calendar year. If patients are shifting utilization by purchasing more medication when it is cheap in the end of the calendar year and less medication when it is expensive in the beginning of the calendar year, then this adherence metric will capture this with a positive coefficient. If this coefficient is substantially larger than the Q4 coefficient, that would be evidence of patients stockpiling medication from Q4 (when it is cheap) to the following Q1 (when it is expensive again).

Tables 2.7 and 2.11 show the results of these adherence regressions for the three dependent variables. As with the previous zero OOP results, the impacts are largest in the 2015 and 2016+ cross-coefficients when the medical spending and pharmaceutical spending were combined in 2015 and the family regulations were adjusted in 2016. I find positive and marginally-significant coefficients for branded and generic medications for calendar year MPR and fourth quarter days filled metrics in 2015. While the coefficients remained positive they were no longer significant for 2016. These coefficients represent modest impacts on adherence. An increase in pre-ACA exposure to the OOP limit by 10% (the rough equivalent of the difference in exposure between middle quintile of products and the highest quintile)

is associated with a 1.2 percentage point increase in calendar year MPR for brands and a 0.9 percentage point increase in calendar year MPR for generic molecules in the 2015 policy regime. The magnitudes are similar but not statistically significant for the post-2015 policy regime. Fourth quarter adherence increases slightly more than proportionally indicating that these gains are at least partially stronger in the end of the calendar year than in other times of the year. However, the fact that most of the adherence gain is not found in the fourth quarter may be some directional evidence that patients are forward-thinking when purchasing products and purchase more earlier knowing that they may reach the OOP limit in the future.

The difference in days of therapy between the fourth quarter and the following first quarter of the calendar year yield positive and significant coefficients for branded and generic products in 2015. This is to say that the impact on adherence in quarter one is smaller than the impact in quarter four when the products are cheaper. Had we observed that the Q4-Q1 coefficients were twice as large as the Q4 coefficients, that would indicate that patients are completely shifting purchasing from Q1 to Q4. This is not what the data shows, which in combination with the twelve-month adherence findings indicates that total adherence is increased rather than shifted due to this policy.

2.6. Conclusion

Section 1302(c)(1) has been an understudied, but impactful provision in the evaluation of the Affordable Care Act. Over the nine years surrounding the change from 2010 to 2018, the overall proportion of branded prescriptions filled by privately-insured patients that are paid at zero OOP cost increased from 8.5% to 20.2%. Furthermore, this has created a calendar year cyclicality of high prescription drug costs in the beginning of the calendar year for many patients and low prescription costs at the end of the year with 27.4% of patients paying zero dollars OOP for branded prescriptions in December. I use variation between pharmaceutical products' patient populations pre-ACA OOP costs to provide evidence that the 2015 and

post-2015 phases of regulation related to section 1302(c)(1) was directly responsible for a portion of this increase.

The increase in the prevalence of OOP limits in private insurance has wide implications for policymakers, industry stakeholders, and the US healthcare system as a whole. Policymakers will continue to monitor the maximum OOP limit to ensure that it does not become so high that it is non-binding and privately-insured individuals are exposed to significant financial risk through healthcare crises. They must also ensure it does not become so low that insurers are no longer able to control for inefficient utilization and moral hazard through OOP costs. If premiums do not rise at the same rate as OOP costs over time then either of these extremes become possible. It may become necessary to provide cost-smoothing for patients around the December to January transition or else risk costly reductions in medical adherence. Insurers will need to continue to search for other methods to control healthcare costs as their ability to control costs via OOP costs has diminished due to this regulation. Pharmacies may consider providing patients with the ability to purchase pharmaceutical products with debt to offset the large changes in prices throughout the plan year. Pharmaceutical manufacturers must understand that their customers face large changes in OOP costs throughout the year and continue to find ways to help their patients maintain adherence to their medications. A part of the proliferation of manufacturer-provided coupons may be a direct response to the maximum OOP limit regulation studied in this paper. Finally, patients must become more sophisticated in their understanding of the insurance contract and may need to respond to price variances throughout the year with increased savings near the end of the calendar year and utilization of manufacturer-sponsored coupons.

The results of this paper prompts research further in many directions. Research related to non-linear pricing in healthcare will become more important than ever as prices become more non-linear throughout the calendar year. This policy will also open the door to increased research in the dynamics of patient response to OOP costs as the plan year start date becomes

a larger discontinuity over which future healthcare prices change as a result of whether a healthcare interaction occurs before or after that date. Additionally, section 1302(c)(1) may have been impactful for hospitals and clinics as well as pharmaceuticals. Medical care is a large proportion of overall healthcare costs so understanding the impact of the regulation in this setting may also be valuable to the health research and policy communities.

TABLE 2.1. Pre-ACA Expected Exposure to OOP Limit

Method:	Drug OOP Only (2014)	Family (2015)	Family-Individual (2016)
100% of Limit	1.0%	4.1%	6.2%
80% of Limit	1.9%	6.2%	9.2%
67% of Limit	3.4%	9.3%	13.2%

Notes: Proportion of patients, weighted by prescription-drug volume and survey weight, who are expected to be exposed to the three different regimes of section 1302(c)(1) based on their spending in the pre-ACA period from 2008-2013. Source: MEPS.

TABLE 2.2. IQVIA Data Summary Statistics

	2012	2018
IQVIA LAAD Prescriptions Data		
Total no. of branded products in panel	734	734
Total sales of branded prescriptions in panel	$74.1B	$108.2B
Total sales of branded prescriptions in LAAD	$75.8B	$175.5B
No. of pharmacies covered in LAAD sample	28,192	48,230
Total no. of patients purchasing brands in panel	57,103,967	50,414,928
Payment Channel		
Private Insurance	54.3%	43.0%
Medicare	26.1%	31.8%
Medicaid	11.1%	19.1%
Other	8.6%	6.1%
Demographics of privately insured using panel products		
% Female	55.5%	56.4%
% Age 0-18	16.4%	14.2%
% Age 19-44	29.1%	30.4%
% Age 45-64	42.4%	44.6%
% Age 65+	12.0%	10.8%

Notes: Summary statistics for balanced panel of branded products in IQVIA LAAD dataset.

2.7. Tables

TABLE 2.3. MEPS Diagnosis Prevalences and Demographics in the Under-65 Privately Insured Cohort by Prescription Utilization

Cohort:	All Privately Insured		with Any Prescription		with Branded Products		with Generic Molecules	
	2012	2018	2012	2018	2012	2018	2012	2018
High Blood Pressure	24.8%	22.4%	32.6%	30.9%	40.1%	36.8%	35.4%	33.8%
High Cholesterol	24.9 %	21.1%	32.1%	28.9%	41.9%	34.9%	34.0%	31.0%
Diabetes	5.7%	4.9%	8.4%	8.1%	14.6%	19.3%	8.7%	8.8%
Arthritis	17.7%	14.2%	23.9%	19.6%	32.8%	29.4%	25.6%	21.3%
Cancer	6.6%	6.1%	8.7%	8.1%	11.7%	10.5%	9.0%	8.4%
% Female	50.8%	50.4%	56.2%	55.5%	56.7%	56.6%	55.5%	54.6%
Age								
0-18	25.2%	24.8%	20.3%	18.6%	15.0%	16.7%	19.7%	17.7%
19-44	40.8%	42.2%	37.6%	39.0%	30.3%	31.9%	35.9%	37.3%
45-64	34.0%	33.0%	42.1%	42.4%	54.7%	51.4%	44.4%	45.0%
% of MEPS	55.8%	58.4%	34.4%	33.0%	13.7%	7.9%	30.0%	28.7%

Notes: Common diagnoses and descriptive statistics from MEPS full-year consolidated 2012 and 2018 files. Diagnosis proportions are from respondents who answered diagnosis questions and are weighted based on survey weight. "All Privately Insured" includes all privately-insured individuals under 65 years of age. "With Any Prescription" is the sub-cohort of those consumers who filled at least one prescription in MEPS while "with Branded Products" and "with Generic Molecules" are the sub-cohorts of consumers who filled products in my branded and generic product panels. A cohort of patients using products which are generic but compete with brands are not shown.

TABLE 2.4. Branded Cross-tabulation by Pre-ACA Exposure to 2016 OOP Maximum Regulation

Exposure Quintile	Mean Exposure	2012 Mean MPR	2012 Mean Q4-Q1 Diff. DoT	2018 Mean MPR	2018 Mean Q4-Q1 Diff. DoT	n	2012 Mean Monthly Cost	% Diabetes Therapy	% Psychotherapeutics	% Neurological	% Antihyperlipidemic
1	0.01	0.65	-4.93	0.75	-1.25	51	294.16	0.06	0.16	0.06	0.00
		(0.04)	(0.71)	(0.04)	(1.00)						
2	0.03	0.64	-5.83	0.72	-3.19	53	217.57	0.19	0.09	0.06	0.02
		(0.04)	(0.76)	(0.04)	(1.09)						
3	0.05	0.74	-3.68	0.85	0.15	53	178.01	0.15	0.06	0.04	0.08
		(0.03)	(0.70)	(0.03)	(0.78)						
4	0.07	0.76	-4.13	0.87	-0.35	50	214.99	0.08	0.18	0.04	0.10
		(0.03)	(0.64)	(0.03)	(0.61)						
5	0.16	0.69	-5.08	0.79	-0.53	54	690.83	0.07	0.06	0.15	0.00
		(0.03)	(0.73)	(0.03)	(0.80)						

Notes: Branded products in the analysis were separated into quintiles based on their exposure to the OOP maximum regulation to observe the characteristics of low compliers compared to high compliers. Source: Exposure data from MEPS. All other fields from IQVIA

TABLE 2.5. The Effect of ACA OOP limit on No Cost Prescriptions: Branded Products

Dependent Variable:	Dec. Zero %	Full-Year Zero %
Drug Spend >100% × Yr = 2014	0.18	0.15
	(0.12)	(0.15)
Total Spend Fam >100% × Yr = 2015	0.41***	0.15
	(0.15)	(0.13)
Total Spend Fam Ind >100% × Yr = 2016+	0.36***	0.28**
	(0.13)	(0.13)
Year and Product Fixed Effects	Yes	Yes
No. of Products	238	238
No. of Years	7	7

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: Two separate estimates of Equation 1 for branded products. Branded products are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.6. The Effect of ACA OOP limit on No Cost Prescriptions: Generic Molecules

Dependent Variable:	Dec. Zero %	Full-Year Zero %
Drug Spend >100% × Yr = 2014	0.02	0.01
	(0.12)	(0.08)
Total Spend Fam >100% × Yr = 2015	0.26***	0.09*
	(0.07)	(0.05)
Total Spend Fam Ind >100% × Yr = 2016+	0.28**	0.06
	(0.12)	(0.09)
Year and Product Fixed Effects	Yes	Yes
No. of Products	318	318
No. of Years	7	7

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Notes: Two separate estimates of Equation 1 for generic products. Generic molecules (all non-brand products with the same active ingredient) are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.7. Adherence Regression: Brands

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	−0.03	1.28	−1.39
	(0.08)	(8.14)	(6.77)
Total Spend Fam >100% × Yr = 2015	0.12*	12.02*	10.62**
	(0.07)	(7.15)	(4.10)
Total Spend Fam Ind >100% × Yr = 2016+	0.10	9.56	5.89
	(0.07)	(7.49)	(4.50)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	224	224	224
Num. obs.	1568	1568	1568

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

All branded products are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.8. Adherence Regression: Brands with pre-ACA MPR less than 50%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	0.12*	4.93	−2.00
	(0.07)	(6.07)	(4.87)
Total Spend Fam >100% × Yr = 2015	0.30***	21.99**	10.96**
	(0.10)	(8.62)	(4.81)
Total Spend Fam Ind >100% × Yr = 2016+	0.35***	28.07***	10.26
	(0.11)	(9.87)	(6.43)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	52	52	52
Num. obs.	364	364	364

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Branded products with 50% or lower initial MPR are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.9. Adherence Regression: Brands with pre-ACA MPR between 50% and 80%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	0.01	13.94	10.20
	(0.19)	(22.31)	(20.97)
Total Spend Fam >100% × Yr = 2015	0.21	25.36	21.13
	(0.13)	(14.83)	(13.94)
Total Spend Fam Ind >100% × Yr = 2016+	0.10	9.19	9.36
	(0.08)	(8.76)	(8.23)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	80	80	80
Num. obs.	560	560	560

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Branded products with 50% to 80% initial MPR are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.10. Adherence Regression: Brands with pre-ACA MPR over 80%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	−0.24	−19.68	−22.05
	(0.17)	(19.07)	(19.99)
Total Spend Fam >100% × Yr = 2015	−0.23	−23.61	−6.43
	(0.16)	(17.76)	(11.07)
Total Spend Fam Ind >100% × Yr = 2016+	−0.13	−6.23	1.05
	(0.12)	(12.40)	(7.36)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	92	92	92
Num. obs.	644	644	644

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Branded products with 80% or higher initial MPR are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.11. Adherence Regression: Generics

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	−0.03	−4.81	−0.29
	(0.06)	(5.98)	(3.42)
Total Spend Fam >100% × Yr = 2015	0.09*	9.27**	3.16**
	(0.05)	(4.46)	(1.23)
Total Spend Fam Ind >100% × Yr = 2016+	0.08	4.44	−0.53
	(0.08)	(5.78)	(2.31)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	257	257	257
Num. obs.	1799	1799	1799

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Generic products are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.12. Adherence Regression: Generics with pre-ACA MPR less than 50%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	0.07	1.73	−6.94
	(0.09)	(8.16)	(5.00)
Total Spend Fam >100% × Yr = 2015	0.11*	8.78	1.82
	(0.06)	(5.36)	(1.41)
Total Spend Fam Ind >100% × Yr = 2016+	0.03	−0.21	−1.02
	(0.11)	(8.87)	(1.88)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	110	110	110
Num. obs.	770	770	770

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Generic products with 50% or lower initial adherence are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.13. Adherence Regression: Generics with pre-ACA MPR between 50% and 80%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	−0.13	−3.86	13.27
	(0.26)	(19.23)	(9.55)
Total Spend Fam >100% × Yr = 2015	−0.04	0.67	3.88
	(0.06)	(5.89)	(2.34)
Total Spend Fam Ind >100% × Yr = 2016+	−0.04	−2.14	1.82
	(0.10)	(8.36)	(2.88)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	71	71	71
Num. obs.	497	497	497

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

Generic products with 50% to 80% initial adherence are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.14. Adherence Regression: Generics with pre-ACA MPR over 80%

	Calendar Yr MPR	Q4 Days Filled	[Y]:[Q4] to [Y+1]:[Q1] Diff. in Days Filled
Drug Spend >100% × Yr = 2014	0.08	3.70	−0.64
	(0.06)	(6.01)	(3.96)
Total Spend Fam >100% × Yr = 2015	0.05**	4.80**	1.20
	(0.02)	(2.31)	(1.38)
Total Spend Fam Ind >100% × Yr = 2016+	0.05	6.71**	3.62**
	(0.03)	(3.28)	(1.75)
Year and Product Fixed Effects	Yes	Yes	Yes
No. of Products	76	76	76
Num. obs.	532	532	532

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$

Generic products with over 80% initial adherence are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded. Cluster robust standard errors are in parentheses.

TABLE 2.15. Branded Products Parallel Trends Test

	Dec. Zero %	Dec. Zero %
Total Spend Fam > 100% X Yr = 2013	0.07	
	(0.05)	
Total Spend Fam Ind > 100% X Yr = 2013		0.04
		(0.04)
Year and Product Fixed Effects	Yes	Yes
No. of Products	275	275
Num. obs.	550	550

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Notes: Output from test of parallel trends in December zero-dollar OOP proportion between 2012 and 2013, based on estimated future exposure to 2015 and 2016 maximum OOP limit regulations. Products are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded.

2.8. Figures

FIGURE 2.1. The Annual Proportion of Private Insurance Branded Prescriptions Paid at Zero OOP Cost

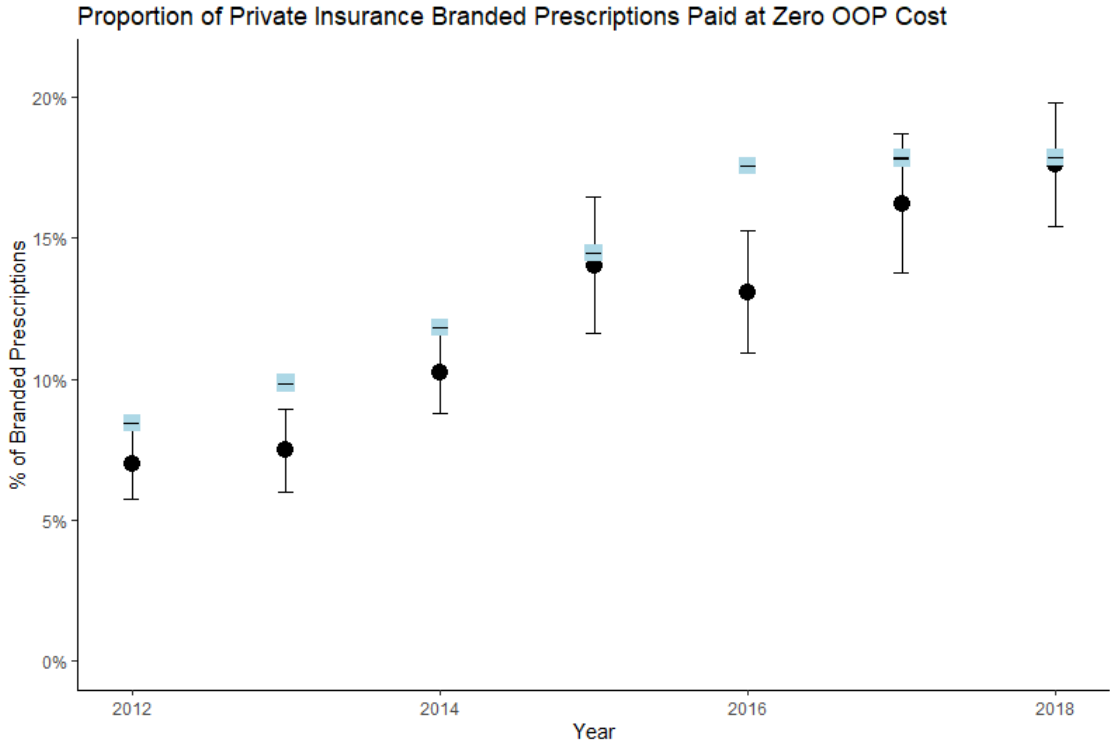

Notes: The proportion of branded prescriptions acquired by individuals with private insurance. MEPS data (black circles) is calculated as the survey-weighted proportion of prescriptions which were purchased for zero dollars OOP. The black error bars represent the 95% confidence interval. IQVIA LAAD data (blue squares) is calculated as the proportion of 30-day-normalized prescriptions which were purchased for zero dollars OOP. Both data series cover the cohort of branded products with over $10 MM revenue in 2012.

FIGURE 2.2. The Annual Proportion of Private Insurance Generic Prescriptions Paid at Zero OOP Cost

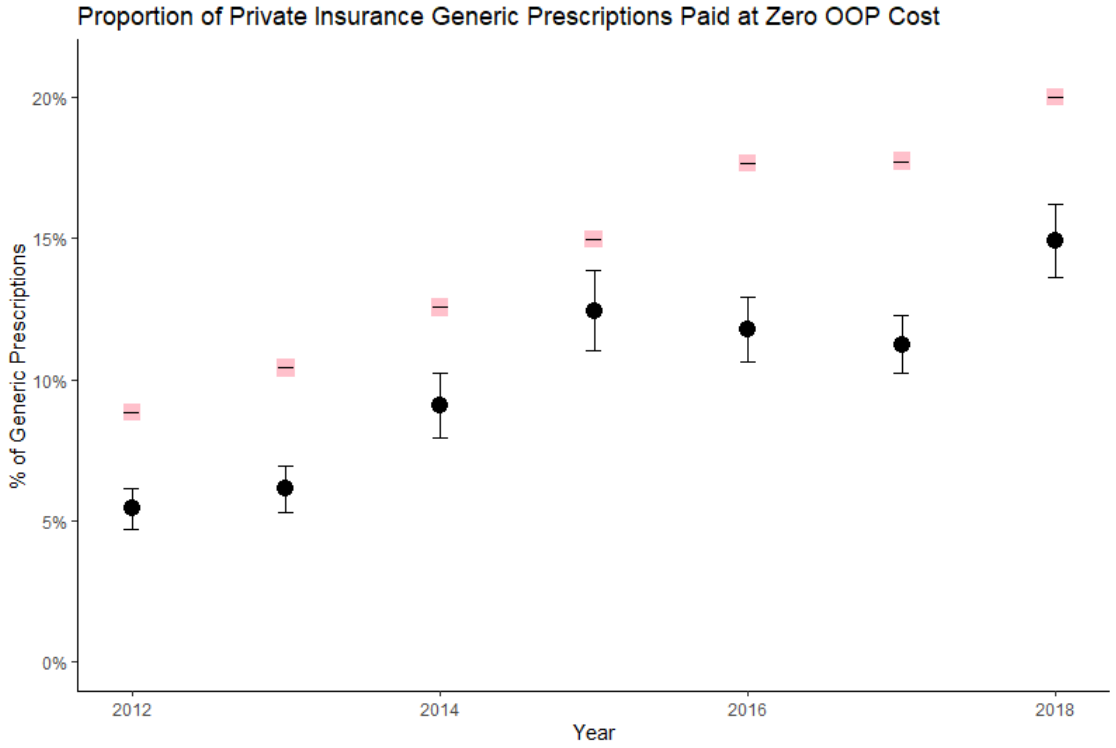

Notes: The proportion of generic prescriptions acquired by individuals with private insurance. MEPS data (black circles) is calculated as the survey-weighted proportion of prescriptions which were purchased for zero dollars OOP. The black error bars represent the 95% confidence interval. IQVIA LAAD data (pink squares) is calculated as the proportion of 30-day-normalized prescriptions which were purchased for zero dollars OOP. Both data series cover the cohort of generic products with over 500,000 days of therapy sold in the 2012 IQVIA dataset.

FIGURE 2.3. The Monthly Proportion of Private Insurance Prescriptions Paid at Zero OOP Cost

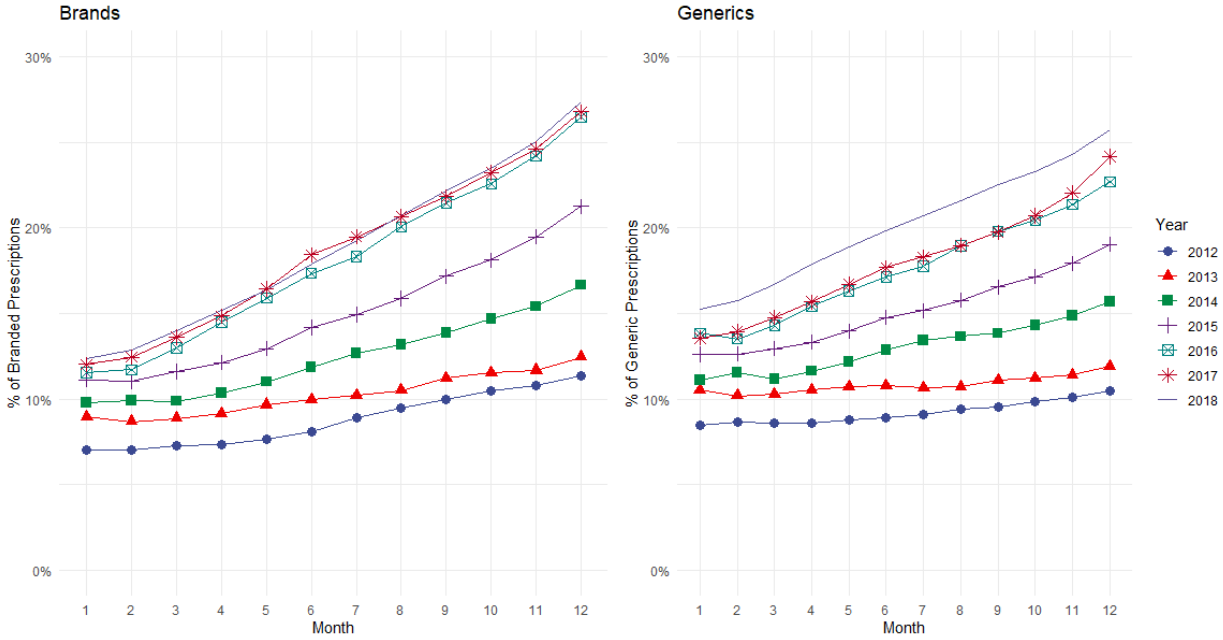

Notes: The proportion of 30-day-normalized prescriptions acquired at zero OOP cost by individuals with private insurance. The cohort of products included are branded products with over $10 MM revenue in 2012 and generic products with over 500,000 days of therapy sold in 2012 IQVIA dataset. Source: IQVIA LAAD data

FIGURE 2.4. Individual OOP Cost Distributions

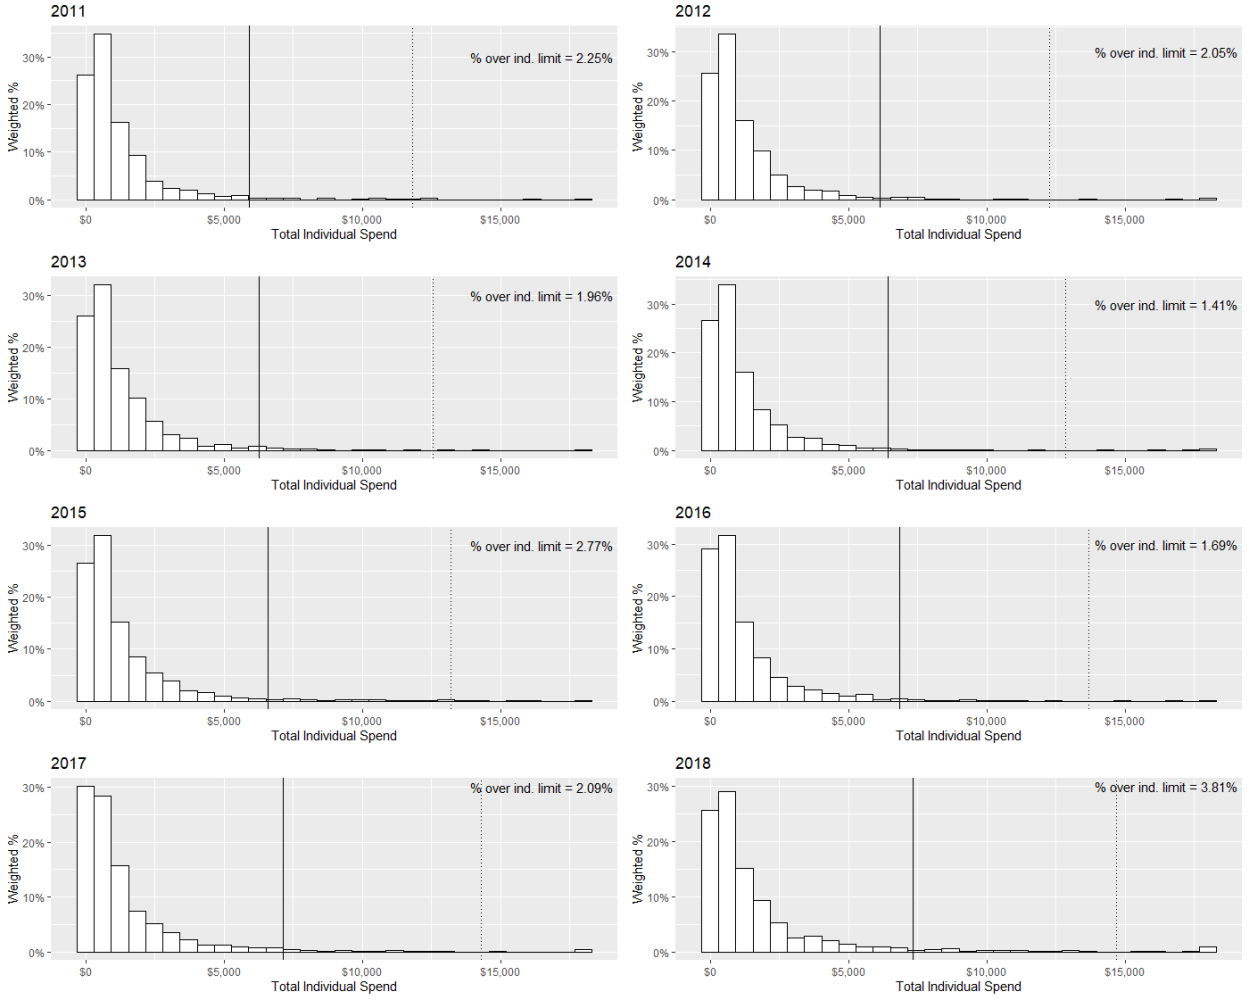

Out-of-pocket healthcare spending distributions for privately insured individuals under age 65 from 2011 to 2018. The individual out-of-pocket limit is shown as a solid vertical line. For 2014 and after, it is the actual limit imposed by HHS. For 2013 and earlier,the 2014 limit adjusted by healthcare CPI. The dotted line is the family out-of-pocket limit for reference. Source: survey-weighted MEPS

FIGURE 2.5. Family OOP Cost Distributions

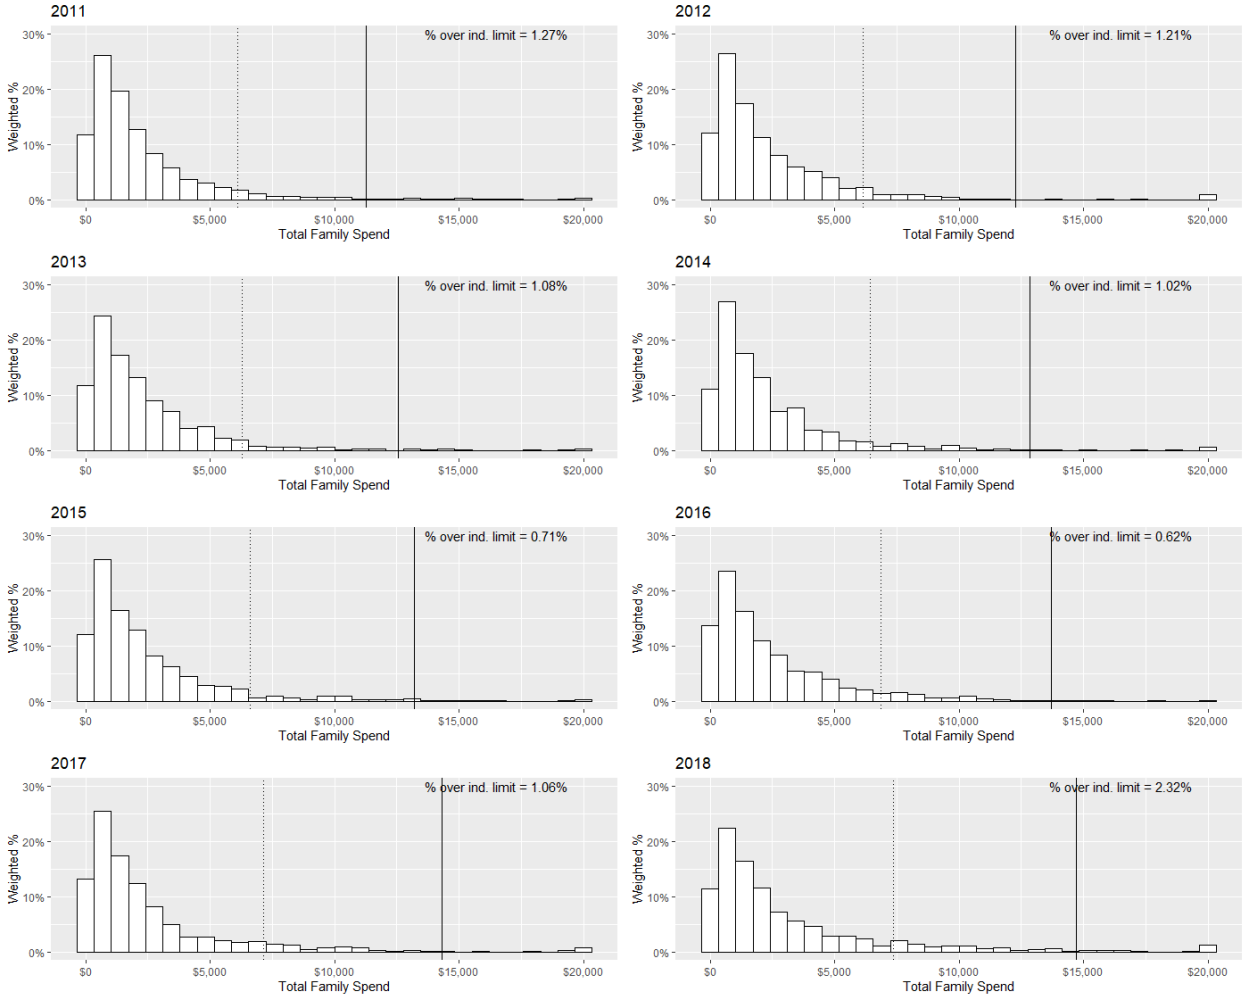

Out-of-pocket healthcare spending distributions for privately insured individuals under age 65 from 2011 to 2018. The family out-of-pocket limit is shown as a solid vertical line. For 2014 and after, it is the actual limit imposed by HHS. For 2013 and earlier, the 2014 limit adjusted by healthcare CPI. The dotted line is the individual out-of-pocket limit for reference. Source: survey-weighted MEPS

FIGURE 2.6. Distribution of Brands by Exposure to 2016 OOP Limit

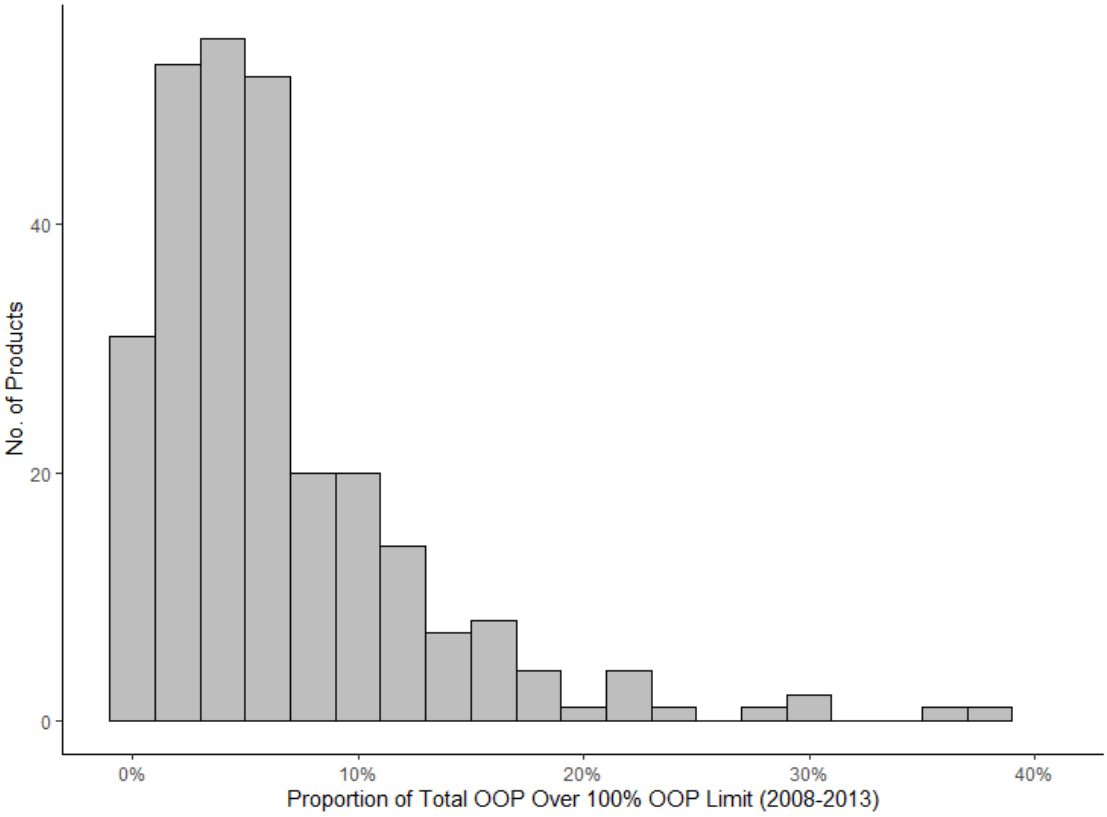

Notes: The horizontal axis represents the proportion of a branded product's patient population that spent over the 2016 OOP limit (in medical-deflated dollars) in the years before the ACA went into effect (2008-2013). Products included are those that had 2012 revenue over $10MM and had a sample of over 30 privately-insured patients in 2008 to 2013. 275 products in total. Source: MEPS

FIGURE 2.7. Difference Between 2015 and 2016 OOP Limit Rules by Brand

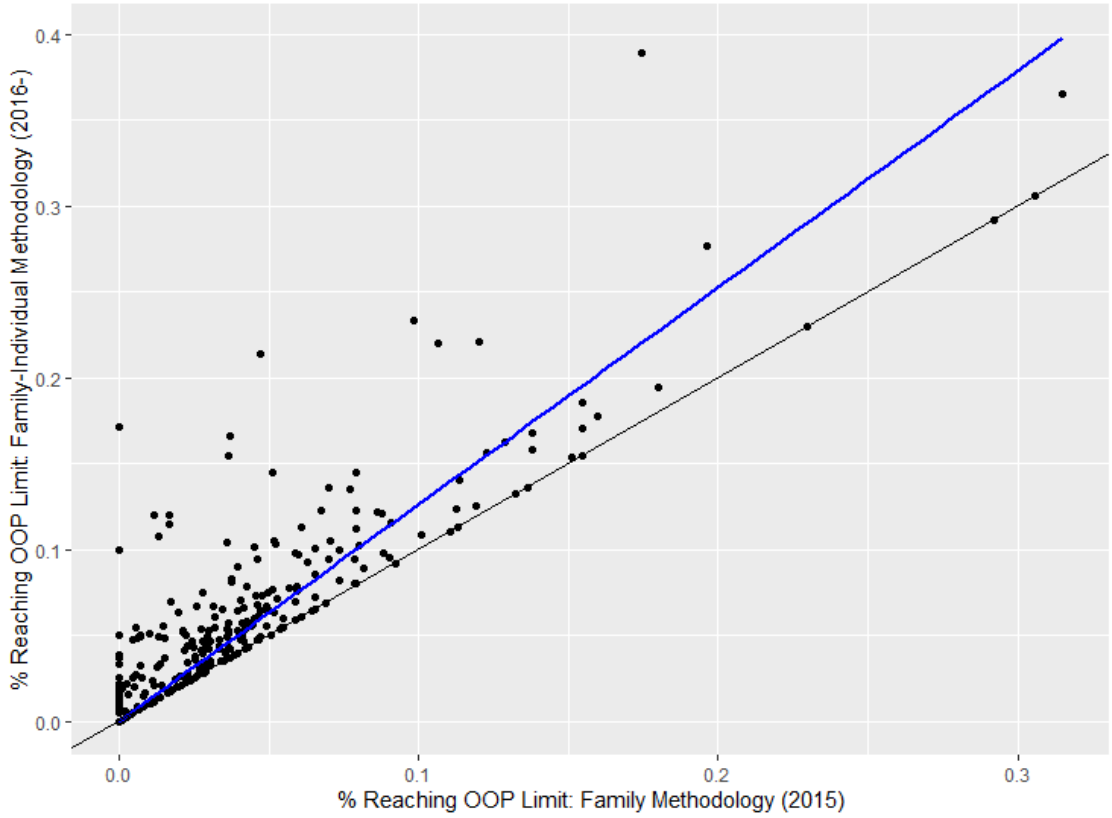

Notes: A scatter plot of branded pharmaceutical products' estimated exposures to the 2015 and post-2015 enforcements of Section 1302(c)(1). The 2015 exposure treats individuals within families as being required to reach the total family limit (which is statutorily twice the individual limit) while the post-2015 enforcement exposure allows individuals within families to reach either limit before the OOP limit is reached. The post-2015 limit is at least as easy to reach for products so no products fall below the 45 degree black line. The blue line is an OLS regression of the data.

FIGURE 2.8. The Effect of ACA OOP limit on No Cost Prescriptions: Monthly Effects Relative to January

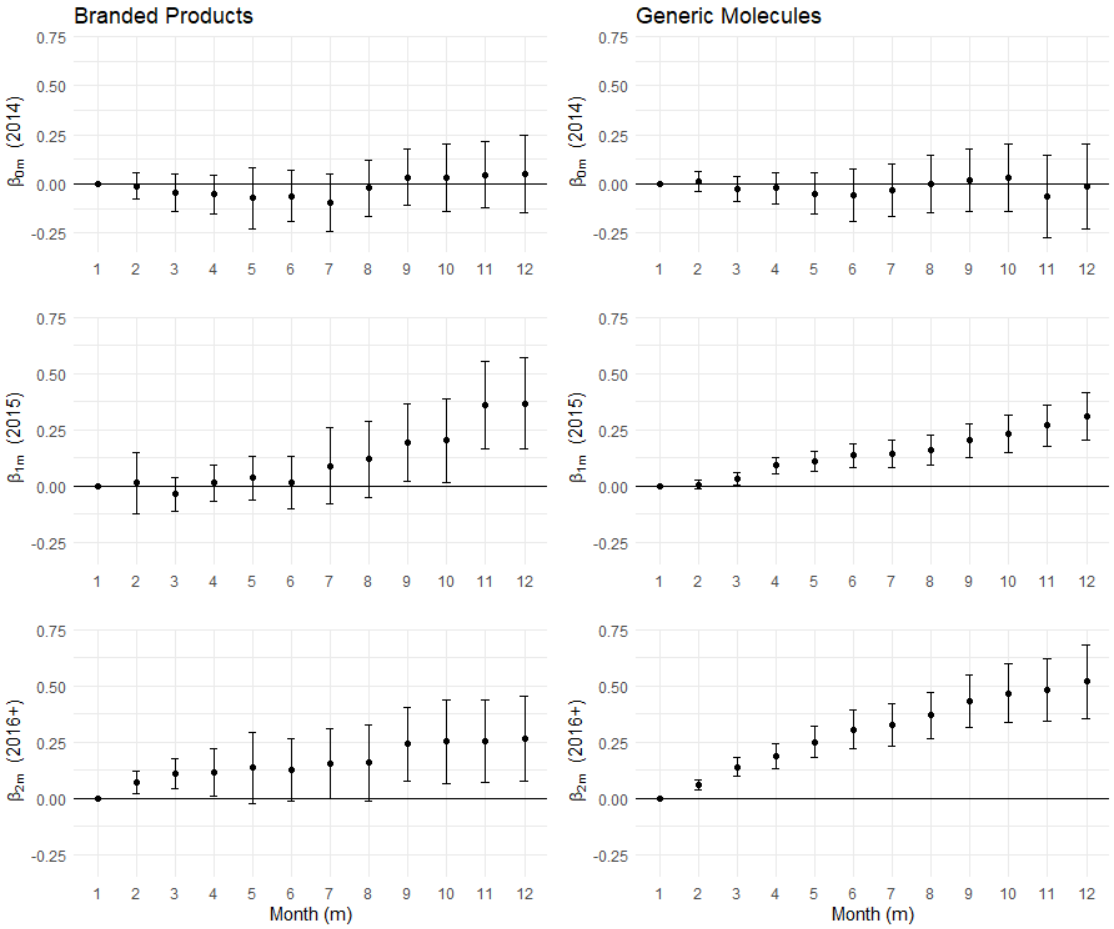

Notes: The chart displays the interaction coefficients from estimates of Equation 2 for branded products (left) and generic molecules (right). Cluster robust standard errors were calculated and 95-percent confidence intervals are shown on the chart. Products are weighted based on their sample in MEPS to de-weight products with high variance in the annual OOP cost distributions. Products with fewer than 30 MEPS patients between 2008 and 2013 are excluded.

FIGURE 2.9. Pre-Trends in Zero-Dollar Prescriptions from 2012 to 2013

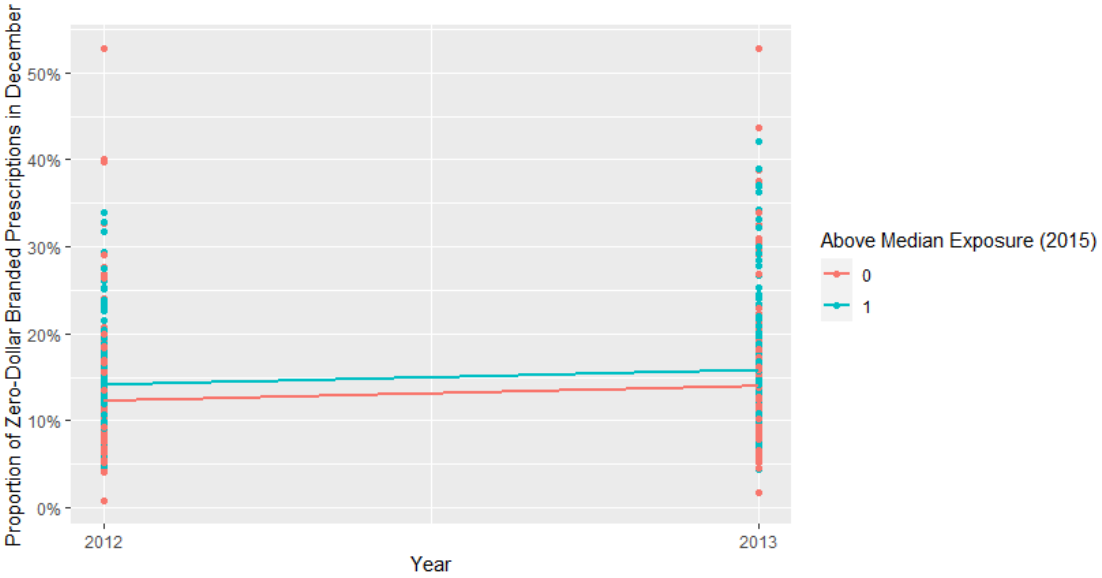

Notes: The proportion of zero-dollar prescriptions in December 2012 and 2013 for branded products included in the regression analysis. The two linear regressions show the change in the average proportion for the products above and below the median exposure to the ACA maximum OOP limit regulation as enforced in 2015.

FIGURE 2.10. Pre-Trends in Zero-Dollar Prescriptions from 2012 to 2013

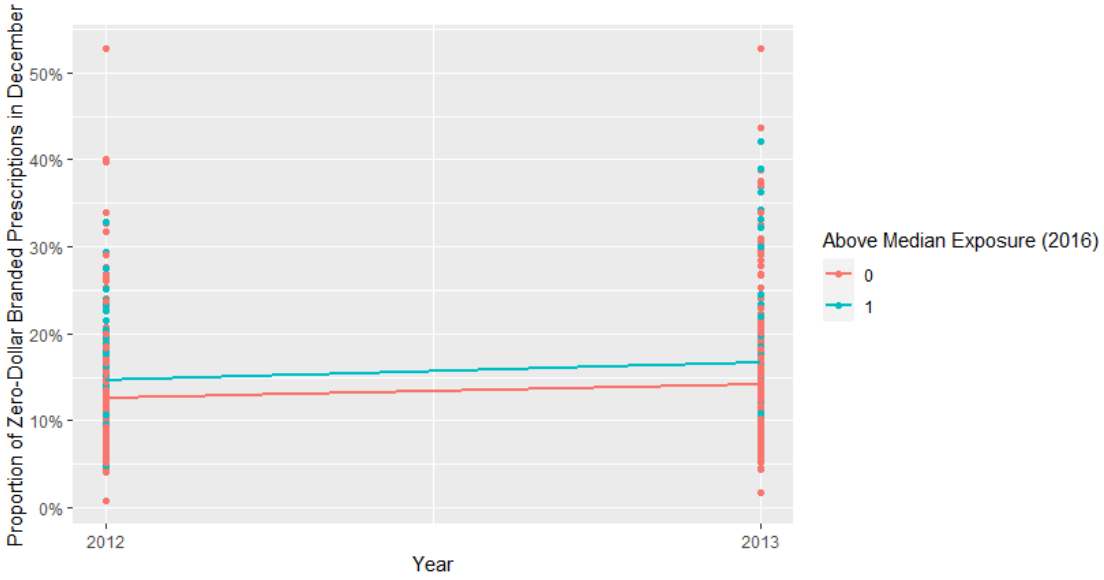

Notes: The proportion of zero-dollar prescriptions in December 2012 and 2013 for branded products included in the regression analysis. The two linear regressions show the change in the average proportion for the products above and below the median exposure to the ACA maximum OOP limit regulation as enforced in 2016.

FIGURE 2.11. An excerpt from the passed ACA legislation: Section 1302(c)(1)

(c) REQUIREMENTS RELATING TO COST-SHARING.—

(1) ANNUAL LIMITATION ON COST-SHARING.—

(A) 2014.—The cost-sharing incurred under a health plan with respect to self-only coverage or coverage other than self-only coverage for a plan year beginning in 2014 shall not exceed the dollar amounts in effect under section 223(c)(2)(A)(ii) of the Internal Revenue Code of 1986 for self-only and family coverage, respectively, for taxable years beginning in 2014.

(B) 2015 AND LATER.—In the case of any plan year beginning in a calendar year after 2014, the limitation under this paragraph shall—

(i) in the case of self-only coverage, be equal to the dollar amount under subparagraph (A) for self-only coverage for plan years beginning in 2014, increased by an amount equal to the product of that amount and the premium adjustment percentage under paragraph (4) for the calendar year; and

(ii) in the case of other coverage, twice the amount in effect under clause (i).

If the amount of any increase under clause (i) is not a multiple of $50, such increase shall be rounded to the next lowest multiple of $50.

FIGURE 2.12. Cumulative Proportion of Patients Reaching OOP Limit

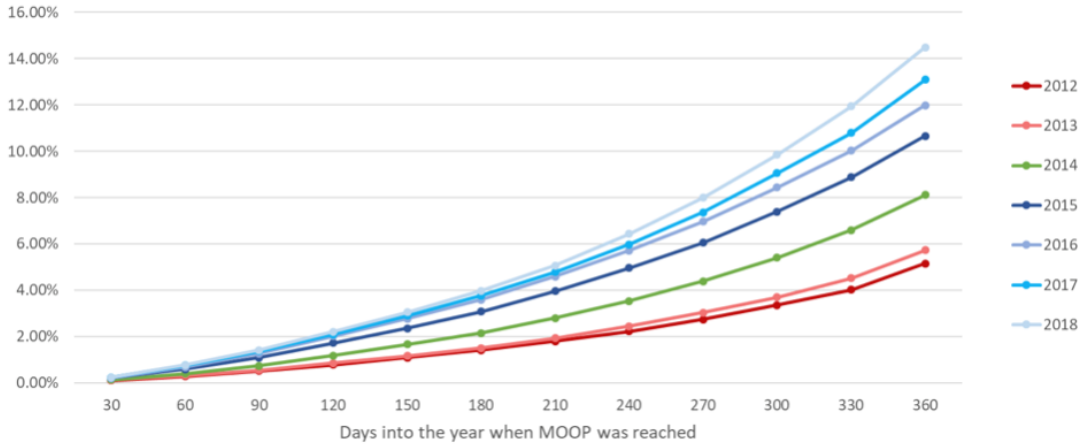

Notes: The cumulative proportion of patients who are classified as having reached the OOP limit a certain number of days into the calendar year based on longitudinal analysis of their prescription claims. Source: IQVIA LAAD.

2.9. Appendix

2.9.1. Parallel Trends for Main Specification. While the available LAAD data only went back through 2012, I used 2012 to 2013 to test for parallel pre-trends between high-exposure and low-exposure products. A regression similar to the main specification performed on those pre-implementation years indicate no significant interaction between the two measures of exposure (corresponding to the 2015 and 2016 and beyond enforcement regimes) and the zero-dollar proportion of branded prescriptions in December of those years. However, as seen in Table 2.15, small positive coefficients are found for both exposure measures. Additionally, when comparing the slopes between cohorts of products above and below median exposures in Figures 2.9 and 2.10, there are no substantial deviations from parallel trends. This provides evidence that there were not unrelated factors before implementation in 2014 that caused products with high exposure to have strongly deviating trends in the proportion of zero-dollar prescriptions to those with low exposure.

2.9.2. Longitudinal Zero-Dollar Flagging Methodology. To further provide evidence that the zero-dollar claims are due to patients reaching an OOP limit, I perform an analysis on patients' prescription claims longitudinally throughout the calendar year to identify the date when patients reach the OOP limit. The previous methods only examined individual claims whereas using IQVIA LAAD data, deidentified patients can be tracked as they fill prescriptions over time. One limitation of the LAAD data is that patients' insurance contract details are not directly available. In other words, there is no information on whether patients' spending is subject to an OOP limit or if that limit has been reached yet. However, I create a panel of patients for each calendar year and track the OOP costs for each prescription to identify a date after which all prescriptions are paid for at zero dollars.

The panel of patients was created requiring each patient to have at least one prescription with full data in each quarter of the calendar year. The IQVIA longitudinal prescription

claims dataset contains information on 90% of retail prescriptions and 60-80% of mail prescriptions in the US (IQVIA Institute 2022), however it does not contain complete information about costs and coordination of benefits for all of these claims since only some data suppliers provide all the necessary fields to understand which plans paid for each claim and how much each paying party paid for a given claim. One concern with creating a panel of patients in this data environment is to exclude patients who have some full data prescriptions and others that have partial or no data. These are patients that use multiple pharmacies, some of which are in the data and others are not fully captured. So requiring valid claims in each quarter reduces the likelihood of including a patient that changed their primary pharmacy mid year from or to a non-captured pharmacy. I also require that over 80% of visible claims in each quarter fully capture costs and coordination of benefits. This reduces the likelihood of including patients who use multiple pharmacies. However, it is still possible that some patients are included that use fully captured pharmacies and fully excluded pharmacies. These patients may create uncertainty in the date of when they reached the OOP limit because there may be more distance between the last paid prescription and the first zero-dollar prescription. These two requirements (at least one valid claim per quarter and over 80% of visible claims with valid cost data in each quarter) lead to a sample of patients that are not representative of all US consumers (many of whom have no prescriptions at all), but are hopefully representative of consumers who have at least one prescription each quarter and who don't oscillate between many pharmacies. This is an important population because these are people who take regular medications and are the primary customers of pharmacies and pharmaceutical companies. It also includes patients who predominately use generic medications unlike the branded-only data shown above.

This panel was examined to find the first day of the calendar year on which all subsequent claims were zero-dollar OOP claims for the remainder of the year. Figure 2.12 shows the cumulative proportion of patients who pay zero OOP for the rest of the year after a given

number of days into the calendar year on the horizontal axis. A trend is visible in this data corresponding to the phases of the maximum OOP limit regulation. Pre-2014 before the major provisions of the ACA went into effect, the proportion reaching an OOP limit was low, then 2014, the transition "safe-harbor year" saw an increases in the proportion of patients reaching their OOP limits in pharmaceutical claims. Then by 2015 when insurers were required to merge pharmacy and medical OOP limits, the proportion increased again. The last regulatory tweak in 2016 affecting individuals on family plans is not particularly visible in this chart although the last four years do show gradual increases in the proportion of patients observed to reach an OOP limit. This data supports the idea that these increasingly common zero-dollar claims fit the pattern of patients reaching OOP limits rather than the alternative hypotheses of workers' compensation or value-based pricing.

CHAPTER 3

A Bad Day for a Heart Attack: Discontinuity in Prices at the Plan-Year Reset

3.1. Introduction

Over the past 20 years, private health insurance plans[1] in the United States have evolved in two seemingly opposite directions: Firstly, deductible[2] phases have become more prevalent and larger: In 2009, 17% of covered workers were enrolled in a plan with an annual deductible higher than $1,000. By 2023, that figure was 57% (Kaiser Family Foundation 2023b). Similarly, the average deductible for health exchange plans was $1,181 in 2014 which nearly tripled to $3,057 in the most recent data (Kaiser Family Foundation 2023a). This means more of the initial risk and cost burden of adverse health events have shifted to consumers over this time period. Secondly, patients are spending more time having reached their out-of-pocket (OOP) limit[3]. From 2012 to 2018, the proportion of pharmaceutical prescriptions paid with no out-of-pocket costs increased from 8% to 20% (See Figure 2.1). Both of these changes are at least partially policy driven: high deductible healthcare plans were given tax incentives in the Medicare Prescription Drug, Improvement, and Modernization Act of 2003 and the Affordable Care Act placed a maximum OOP limit on non-grandfathered plans in 2014.

[1]Including employer-provided, health-exchange, and individual plans

[2]Deductibles are the amount an insured consumer must pay for covered health care services before their insurance plan starts to pay. With a $2,000 deductible, for example, the consumer must pay the first $2,000 of covered services. Deductibles typically reset every year.

[3]Out-of-pocket limits set a maximum amount of out-of-pocket costs (not including premium payments or out-of-network care) that a consumer must bear within one year. After a consumer pays this amount in one plan-year, all subsequent covered in-network care is paid fully by the insurer

As a result, consumers are more likely to pay full costs for healthcare at the beginning of their plan-year, which often coincides with the beginning of the calendar year, due to the increased prevalence of deductibles, but they are also more likely to pay nothing for in-network healthcare at the end of their plan-year due to the increased prevalence of out-of-pocket limits. These two factors are leading to uneven costs throughout the plan-year for many privately insured consumers who face a large spike in costs on January 1st (potentially increasing from zero cost to full cost). Furthermore, the cost dynamics are different for a consumer who has an expensive healthcare event at the beginning of their plan-year compared to a consumer who has the same healthcare event at the end of their plan-year. The second consumer will almost certainly have to pay through their deductible (if they have one) which has just reset, but they will be able to accrue all that paid cost to their fresh plan-year and their insurance will be likely to pay for a higher proportion of their costs going forward that year. The first consumer's experience varies based on their prior costs: if they had no costs in plan-year prior to the event, they will have to pay their deductible, but they will accrue these costs to their outgoing plan year which will reset and they will have to pay for the next year's deductible as well. If they had many prior costs in the plan-year of the expensive healthcare event, then they will pay relatively little for the event itself because they will likely be in the coinsurance phase or out-of-pocket limit phase of their coverage. But they too will accrue any costs to the outgoing plan-year, and their deductible and OOP limit will reset requiring them to pay through them again in the following year.

In this paper, I utilize a regression discontinuity design centered around the January 1st plan-year reset to measure these ex-post inequities caused by this highly non-linear pricing environment for privately-insured patients who visit the emergency room (ER) with a myocardial infarction (also known as a "heart attack"). I also leverage the discontinuity to measure the elasticity of demand for pharmaceutical products among myocardial infarction

patients. This work can provide insights to questions that are fundamental to developing private and publicly-funded health insurance contracts that are equitable and efficient: Would a 12-month rolling deductible and OOP limit, if it were feasible, be more fair and reduce annual price shocks? Should insurance contracts include OOP price reductions for highly cost-effective medications after an adverse health event like a myocardial infarction?

This analysis reveals several findings: (1) The plan-year discontinuity between those whose myocardial infarction ER visit was in early calendar year compared to late calendar year leads to significant decreases in the average pharmaceutical coinsurance rate, calculated as the amount the patient pays divided by the total amount paid to the pharmacy, of 5.4 percentage points over the first 365 days after a myocardial infarction ER visit, which is equivalent to a 21% relative decrease. This decrease in average pharmacy coinsurance rate is concentrated in the months in the plan-year following the discontinuity where average coinsurance rates decrease between 10 and 12 p.p. (2) The decrease in cost is related approximately equally to an increase in zero-dollar out-of-pocket prescriptions filled and a decrease in full-cost out-of-pocket prescriptions filled. (3) The decrease in cost translates to increases in pharmaceutical expenditure over the year with a point estimate is an increase in total pharmaceutical expenditure of $198 (5.8% relative increase). This results in an estimate of the elasticity of demand for pharmaceutical products of -0.23 which is in line with estimates for overall outpatient healthcare from the RAND health experiment (Manning et al. 1987). (4) This increase in pharmaceutical expenditure is concentrated in branded pharmaceutical products and those that are not related to the circulatory system. No changes in quantity are measurable for circulatory products.

3.2. Background: Measuring Moral Hazard in Healthcare

Over the past decades, health economists and policy researchers have expended tremendous resources and effort to understand how healthcare consumers in the U.S. respond to OOP costs for healthcare treatments. Prominent examples of controlled experiments include

the RAND Health Insurance Experiment between 1974 and 1981 (Newhouse 1993) and the Oregon Health Plan (Finkelstein et al. 2012). Many other researchers have focused on the interface between exiting the private insurance market and entering Medicare at age 65 (Card et al. 2009) (Aron-Dine et al. 2015) (Chandra et al. 2024). Others have collected data from private firms to investigate how those employees respond to changes in healthcare prices (Huskamp et al. 2003) (Goldman et al. 2004) (Brot-Goldberg et al. 2017). Several literature reviews, for example, (Einav and Finkelstein 2018) and (Cutler and Zeckhauser 2000) provide summaries of much of this research. However, all of the various approaches have drawbacks, so it is important for researchers to have multiple approaches to answer the fundamental questions about consumer response to healthcare prices. This study, building off of (Hettinger 2023) and (Kowalski 2016), provides a template for a new approach leveraging discontinuities in prices due to deductible and out-of-pocket limit phases to measure outcomes in time relative to the occurrence of a major health event: in this case, myocardial infarctions.

The first major drawback of many of the above studies is that they are not easily replicable in the current healthcare market landscape. Health insurance experiments are extremely costly and many of the Medicare studies with identification strategies relying on the Part D coverage gap are no longer possible since the coverage gap has been closed. It is an open question whether healthcare consumers' response to high OOP costs has changed over time, but there are theoretical reasons to believe consumer responses might change over time: Firstly, the basket of commonly-used healthcare products changes dramatically on the time scale of decades: many of the most commonly-used therapies today did not exist when the RAND experiment took place. Secondly, many market participants jointly determine which healthcare treatments are delivered to a consumer: doctors, insurers, and medical administration policies all contribute to make medical decisions. If those processes change over time, then so too may the effects of OOP costs. For instance, if doctors, physician

assistants, and nurses have an improved understanding the costs and benefits of alternative treatments, then we should expect consumers' responses to OOP cost to be less helpful in optimizing health outcomes. Alternatively, if healthcare providers become more beholden to the consumers' desires through the ubiquity of physician reviews and insurers become better at determining value, then OOP costs that are proportion to total costs may be more helpful for rejecting low-value treatments. Finally, it is possible that consumers' decision-making processes may change in unexpected ways as society changes as a whole; healthcare consumption must be balanced with other consumption, and health information availability may also change over the decades.

The second major drawback of some of the previous approaches is the limited populations whose price-sensitivity is measurable. The Medicare studies may not be extendable to the majority of the population who are far from the Medicare age threshold. Due to sample size constraints, the experimental and firm-based approaches struggle to measure heterogenous impacts among smaller subsets of consumers who might behave differently from the average employee. Answering the price sensitivity questions for important sub-populations, like those who have experienced a heart attack, is valuable to better understand the mechanisms that affect the elasticity of demand for healthcare. Furthermore, since moral hazard is one of the key market failures that precipitates patient copayment, understanding more precisely *which* parts of healthcare have low or high elasticity of demand could improve optimal insurance contracts.

Several papers[4] have used this injury/plan-reset identification strategy in the past, however this research differs from them in several ways: Firstly, in contrast to (Hettinger 2023), all effects are measured relative to the date of the injury rather than relative to the calendar year. This is an important difference because one would expect medical events (e.g. follow-up visits, prescription refills, or complications) to not be consistent in time relative to the calendar year, but relative to the date of the severe injury. For instance, if we compare 2022

[4]At least, (Card et al. 2009), (Kowalski 2016), and (Hettinger 2023)

calendar year costs of a person who had their myocardial infarction 10 days before the end of 2022 to one who had it 10 days into 2022, the latter would have much higher costs because the ER visit, the hospital stay, and the post-hospital pharmaceuticals will all be purchased in 2022 rather than in 2021. So to compare the two patients across a discontinuity in a similar manner, they should be examined relative to the ER visit. Secondly, this research is performed on data collected after the implementation of the ACA which increased the prevalence of out-of-pocket maximums (Illich 2024) and changed many other aspects of the healthcare system which may affect the magnitude of the cost differential and elasticity of demand. Thirdly, while this analysis only measures pharmaceutical spending due to data limitations on medical spending, it is the only one to capture the effect of post-coupon prices, which are the prices to which patients are actually exposed. Finally, this is the only paper to focus specifically on myocardial infarction patients who may be an important group to understand due to the severity of the condition and the effectiveness of pharmaceutical treatment in preventing future cardiovascular events: Their elasticity of demand for healthcare may be different than the general population.

3.3. Data

The IQVIA Longitudinal Access and Adjudication Data (LAAD) anonymized pharmacy claims dataset contains claim-level information on a substantial proportion of pharmacy transactions in the United States. The discontinuities on January 1st of 2017, 2018, 2019, and 2022 were selected for analysis. These discontinuities were selected as the most recently available data that was not severely affected by the COVID-19 pandemic which began in early 2020. LAAD also contains a medical component that collects information on a smaller proportion of medical claims in the United States and includes procedures performed and diagnoses recorded in emergency rooms.

There are several features of this data that allows for the regression discontinuity approach: Firstly all claims can be attributed to an anonymized patient identifier that allows

me to track patients' claims over the course of several years. This identifier can link patients across medical and pharmacy providers. Secondly, pharmacy and medical claims can be attributed to a primary payer (often an insurer or pharmacy benefit manager) to identify the channel of payment (e.g. private insurance, Medicare, Medicaid). Thirdly, pharmacy claims also contain information about secondary payers, if one exists. These secondary payers are often manufacturer-provided coupons or supplemental insurance. This allows me to observe final prices that patients actually pay after coupons and is information that may not be available in data originating from insurers. Finally, the data contains information related to each claim. On the medical side, this includes HCPCS codes and modifier codes for procedures, ICD-10 codes for diagnoses, service and billing dates, and price information for a portion of the data. On the pharmacy side, this includes the product dispensed, the total reimbursement, the OOP costs (before and after secondary payers), and the number of days over which the prescriber intended the prescription to be used.

To identify which medical claims are emergency room claims, the medical claims with the HCPCS codes 99283, 99284, and 99285 are included in the analysis. Table 3.1 summarizes the distribution of those codes showing that, as expected, myocardial infarctions are very often of the highest complexity. Myocardial infarctions are classified as such based on the ICD-10 diagnosis codes. To ensure that emergency room visits were "surprising" to patients (and therefore unplanned), any myocardial infarction for which a patient had another previous ER visit within 90 days were excluded from the analysis. Table 3.1 also contains summary data for the ICD-10 diagnosis code descriptions used.

Table 3.2 shows the demographic information for patients visiting the ER with a myocardial infarction diagnosis. Males make up 65% of the myocardial infarction ER visits throughout the sample period. Meanwhile 90% of myocardial infarction commercially-insured ER visits are older than age 45. For age and gender, there is very little difference between the

left- and right-sides of the January 1st cut-point. "Placebo" tests for continuity in these variables are performed at the end of Section 3.4 and in Figure 3.5.

3.4. Empirical Strategy

To isolate the effects of a unexpected and idiosyncratic change in prices for privately-insured consumers in the United States, I use a regression discontinuity design (Lee and Lemieux 2010) focused around the change in prices due to an unexpected ER visit occurring near the insurance plan-year reset date. The vast majority of commercial insurance plan-years reset at the beginning of the calendar year on January 1st.[5] On this date, both deductibles and out-of-pocket limits are reset regardless of the prior year spending. Many private insurance contracts have deductibles and all non-grandfathered private health insurance contracts have out-of-pocket limits as of 2015 due to Section 1302(c) of the ACA. Because of this reset, the out-of-pocket spending due to a costly medical event that occurs at the end of December accrues to the plan-year that is ending on December 31 giving the consumer very little time to adjust behavior to the decreased costs. Meanwhile, a medical event occurring in January will accrue to the new plan-year and consumers are more likely to transition through the three phases of coverage: deductible, coinsurance, and OOP limit. Each phase has a decreasing expected out-of-pocket cost from full cost in the deductible phase, to partial cost during the coinsurance phase, to zero cost in the OOP limit phase.

In the context of this study, a regression discontinuity approach attempts to estimate the effects of a heart attack occurring on the last day of the plan-year compared to the first day of the following plan-year. In general, consumers will have lower future costs if the event occurs on the first day (January 1st) of the plan-year rather than on the last day (December 31st)

[5]Kaiser Family Foundation's Health Benefits Annual Survey reports that 69% of covered workers by firms offering health benefits have their plan year begin in January (Kaiser Family Foundation 2017b). However, my analysis of the IQVIA LAAD data from 2014 to 2018 at the plan level showed that out of all plans for which I could positively identify a month when the proportion of zero-dollar claims decreased, 87% to 94% (depending on the year) of the volume occurred in January. It is possible that some plan-years start in other months, but the OOP limits and deductibles reset in January.

because the associated costs due to the heart attack will accrue to the deductible and OOP limit. However these patients on either side of the plan-year reset are otherwise similar as long as there is no manipulation of event date. One concerning form of manipulation would be if consumers who have a heart attack in late December delay treatment until January 1st with the understanding that this will lower their future costs. To limit the possibility of this manipulation, this study focuses on a medical event that often begins in the emergency room and is often severe enough to limit the likelihood of delay. There are many candidate medical events, but this paper focuses on myocardial infarctions, which are emergent and life-threatening, especially without immediate medical care. Even so, I remove December 31st and January 1st from my sample to reduce the likelihood of intentionally delayed treatment on New Year's Eve into New Year's Day morning. Figures 3.1 shows the density of ER visits across the discontinuity including December 31st and January 1st. Myocardial infarction ER visits are consistent throughout the entire data range. Additionally, a McCrary test (McCrary 2008) was performed to test for possible manipulation in ER visit timing, and the test cannot reject the null hypothesis with a p-value of 0.54. Finally, patients older than age 65 were excluded from the analysis due to the likelihood that they are likely to be partially covered by Medicare and their employer which can lead to costs for a myocardial infarction not accruing to their pharmacy benefits due to them having multiple plans.

In the spirit of a regression discontinuity in time approach, all outcomes evaluated in this analysis are measured relative to the injury date. This is critical because the serious injuries which provide the natural experiment in expected pharmaceutical costs are life-changing events that affect future medical decisions. For example, a myocardial infarction often leads to new and more aggressive medical interventions. I assume that these medical interventions related to the myocardial infarction vary in time relative to the initial ER visit. In other words, the assumption is that a patient who enters the ER on December 10th has a similar medical experience on December 11th as a patient who enters the ER on January

10^{th} has on January 11^{th} and this similarity will continue until one of the patients crosses the following plan-year's threshold for the following year. See Figure 3.2 for an illustration of three potential private insurance contracts. For patients near the cut-point, this is not a major issue since there are very few days where one patient has crossed the threshold of the following plan year while the other has not. But for patients farther from the cut-point, this factor is more important. To account for this, most regressions are performed with dependent variables in thirty-day increments relative to the injury date or for the full 365-days which controls for variations in medical behavior throughout the calendar year.

It is important to remember that the difference between these populations on either side of midnight on January 1^{st} is the plan-year to which these medical costs accrue. This difference in cost accrual has two effects for consumers whose heart attack is after the plan-year reset: decreasing the average future OOP prices by moving consumers through the kinks in the insurance contract (Figure 3.2), and increasing the expected OOP prices of the heart attack ER visit due to being more likely to be in the deductible phase of the new plan-year. So the first effect is a future price decrease which has a positive income effect, and the second effect is a current price increase which is a negative income effect. Furthermore, the prices are not easily observable by the consumer until they receive the ER bill. Medical bills are notorious for being delayed relative to the event, and my data supports this with a mean delay between ER visit and ER bill of 48 days (see Figure 3.3 for the distribution of bill wait). Furthermore, some unsophisticated patients may not fully understand their insurance contracts or the extent to which medical payments affect their pharmacy OOP prices. However, consumers are able to observe their pharmacy OOP prices when they purchase their medicines.

I implement this regression discontinuity design using both global parametric regressions as well as local-linear nonparametric regressions.

The following equation is used for the local-linear nonparametric regressions:

$$(1) \qquad Y_i = \alpha + \beta_0 T_i + \beta_1 [InjuryDate]_i + \beta_2 [InjuryDate]_i \cdot T_i + \gamma \mathbf{X_i} + \varepsilon_\mathbf{i}$$

where Y_i is a dependent variable that may be affected by the different cost structures related to having accrued injury-related expenses to the deductible and OOP limit. The primary dependent variables of interest are average coinsurance rates which are calculated as the total pharmaceutical OOP cost to patients (after all coupons and third-party payers) divided by the total pharmaceutical expenditure reimbursed to the pharmacy. The latter quantity does not discount any rebates to insurers which are not publicly available, but it does include dispensing fees charged by the pharmacy. T_i is the treatment dummy variable with "treated" indicating that the injury occurred after January 1$^{\text{st}}$ and was eligible to be accrued to the deductible and OOP limit. $[InjuryDate]_i$ is the number of days that the injury occurred after midnight on New Year's Eve (with negative values indicating that the injury occurred before midnight on New Year's Eve). $\mathbf{X_i}$ is a vector of control variables such as patient age and sex, and pre-injury pharmaceutical utilization, as well as dummy variables to indicate the plan-year[6] The coefficient of interest, β_0 is the discontinuity in the outcome variable between individuals whose ER visits occur after the plan-year cut-point versus before the plan-year cut-point.

The optimally-selected bandwidths for nonparametric regressions are chosen to minimize mean-squared error using the methods described in (Calonico et al. 2020). A triangle kernel is used for the weighting of observations near the cut-point and the bandwidths are flexible to be different widths on either side of the cut-point. The main benefits of local nonparametric approaches are: (1) The concerns related to misspecification of the running variable functional form are alleviated. (2) There are standard econometric methods used to select

[6]Control variables that are not currently available in this research, but could be available in insurance-based data sets include: details about the patient's plan cost structures (e.g. deductible, OOP limit, coinsurance rates), number of family members, and spending prior to the injury.

bandwidths that optimize certain statistical properties. (3) The methods are standardized which prevents researchers from making arbitrary decisions.

However, there are some downsides to local nonparametric regression in this case. (1) The theory underlying the method requires that the running variable is continuous. In my case, time is continuous, but my measurement of the timing of ER visit is not since the data is only available by day. Furthermore, there is uncertainty in the billing procedures around the cut-point. It is unclear how an insurer handle the claims of a person who entered the ER at 11pm and left at 1am: Either they pick a day (which is hopefully the day that is given in the IQVIA medical claims data) or they split the procedures in some other way. Whatever they do, there is fundamentally some uncertainty as to whether the costs accrue in the previous plan-year or the current plan-year. Additionally the costs do not all occur on the day of the ER visit. While the majority of OOP costs occur on the day of the ER visit, OOP costs are substantially elevated for the 3 days following the ER visit (See Figure 3.4). (2) Local nonparametric estimation does not use all the information in creating the estimate. If data far from the cut-point can provide information about the data near the cut-point then this information is left unexploited in local nonparametric regression. In this setting, there is high variance in coinsurance rates and expenditures, and due to the random nature of myocardial infarction occurrence throughout the year, the functional form related to the running-variable is very close to linear so data far from the cut-point can inform the data near the cut-point. (3) While the identifying assumption in regression discontinuity designs is "As-If Random" variation at the cut-point, the case of myocardial infarctions satisfies not only that assumption, but also "As-If Random" variation throughout the whole year. The plausibility of this full-year randomness assumption is supported by the flat slopes of the placebo regression discontinuity tests.

The following equation is used for the global parametric regressions:

$$Y_i = \alpha + \beta_0 T_i + f([InjuryDate]_i) + \gamma \mathbf{X}_i + \varepsilon_i \tag{2}$$

where $f(\cdot)$ is a polynomial that is separated at the cut-point and selected that minimizes the Bayesian information criterion (Hausman and Rapson 2018). The variables are the same as those in the nonparametric specifications and the bandwidth is the full data available between July 8$^{\text{th}}$[7] and June 30$^{\text{th}}$. Finally, I use cluster-robust errors clustered at the mass points caused by daily data as suggested by (Lee and Card 2008). Due to the above methodological considerations, I focus my findings on the global parametric regression results, while reporting the local nonparametric results for robustness.

Figure 3.5 presents a non-outcome regression discontinuity test on the age and sex variables to ensure that the populations on both sides of the cut-point are similar in observable characteristics. The myocardial infarction age and sex variables are consistent throughout the whole gamut of injury dates relative to the cut-point. Patients are also similar in their pharmaceutical spend (including OOP costs and insurance payments) and prescriptions filled in the 90 days prior to their myocardial infarction and ER visit. All p-values are not significant at either the 5% or 10% levels.

3.5. Results

3.5.1. Full-Year Results. One underlying hypothesis this study is testing is whether or not the insurance contract structures among privately insured consumers in the United States lead to differentials in cost-sharing based on the date of an unexpected health crisis. To measure this, I calculate the average coinsurance rate for pharmaceutical products as described in Section 3.4 for the period of 365 days following the first ER visit related to a myocardial infarction[8]. The discontinuity data are shown in Figure 3.6 on the top panel. The linear trend lines shown are the estimated global parametric regression results from Equation

[7]The first week of July is omitted to avoid any potential inconsistencies in ER care due to the "July Effect" when residencies begin in the United States. Conveniently, this is quite distant from the regression discontinuity cut-point for this study of January 1.

[8]The reason for this length of time rather than a shorter or longer time period is so that all individuals, regardless of ER visit date, are being examined for the same exact calendar days with the difference being the order in which they occur relative to the ER vist date.

2 with BIC minimizing separate polynomials (linear, in this case). There is a visible decrease in average coinsurance rate at the January 1st cut-point. The global parametric regression estimate in Table 3.3 show a 5.4 p.p. decrease (95% CI: -4.5 p.p. to -6.3 p.p.) in average pharmaceutical coinsurance rate for those whose ER visit is after the cut-point relative to those before the cut-point. This difference is approximately equivalent to the difference in average coinsurance rates between the 95% coinsurance and 50% coinsurance plans in the RAND Health Experiment where the difference is 6 p.p. (Manning et al. 1987).

So the underlying hypothesis that pharmaceutical cost structures change based on when a health crisis occurs within the plan-year is supported by the data. Further information can be gathered from the slope of the regression in the top panel of Figure 3.6 noting that the average coinsurance rate increases approximately linearly through the course of the calendar year on both sides of the cut-point. The practical explanation for this is that as patients' ER visits are shifted later in the plan-year they have linearly less time to use the better coverage phase resulting from the cost accrual related to the ER visit and more time in the following plan-year where costs are not affected by the accrual from the ER visit.

The main question this study attempts to answer is whether myocardial infarction patients change pharmaceutical consumption due to differentials in pharmaceutical prices. The bottom panel of Figure 3.6 shows the evidence in terms of total pharmaceutical expenditures (OOP plus insurer pay) in the 365 days following the ER visit. The global parametric regression estimate in Table 3.4 shows an increase in pharmaceutical expenditures of $198 (95% CI: $26 to $370). This point estimate in combination with the decrease in average coinsurance rates yields a naïve estimate of the elasticity of demand for pharmaceutical care after a myocardial infarction ER visit of -0.23[9]. This point estimate is consistent with the

[9]This estimate is "naïve" because it does not account for the dynamic nature of pricing where future prices depend on current consumption. It also only captures the prices of products that are actually purchased rather than all prices that patients are exposed to. Future research could partially account for this by evaluating the price on all prescriptions that are presented at the pharmacy rather than only those that are filled.

RAND population elasticity of demand for healthcare estimates which range between -0.1 and -0.4 depending on the method and object of interest.[10] This data can reject an elasticity of demand for pharmaceutical products greater in magnitude than -0.43 for the year following a myocardial infarction ER visit.

As other researchers have pointed out, dynamic incentives matter in health insurance (Aron-Dine et al. 2015). Consumers are forward-looking and consider future prices as well as the spot prices, and the implications of that on these findings must be considered. These dynamics affect the interpretation of the above findings. Those whose myocardial infarctions occur after the plan-year reset are likely to be in lower OOP cost phases of coverage than those whose ER visit was at the end of the previous plan-year. If we assume that those with the lower costs will seek more care during the current plan-year and delay less care into the following plan-year due to the dynamic incentives, then that would lead my estimates of elasticity of demand to be overestimated in magnitude.

One of the most concerning aspects of moral hazard for pharmaceutical products is if behavioral effects combine with high OOP prices limit the purchasing of cost-effective medications leading to detrimental discontinuities in care (Chandra et al. 2024). Another area of concern is whether low OOP prices affect opioid usage which may lead to potential for abuse. To examine these questions in my data, I examine the quantity responses of the discontinuity for pharmaceutical products in four groups: (1) all products, (2) products related to circulatory conditions[11] (3) products not related to circulatory conditions, and (4) opioid medications used for pain relief. Figure 3.7 shows the regression discontinuity charts for all four groups. There is no statistically significant discontinuity measured for all products, circulatory products, or opioid products, but there is a statistically significant

[10]However, I was unable to find any analysis of the RAND study that calculated elasticities for pharmaceutical care directly. This is likely due to the increased importance and cost of pharmaceuticals as a means of care in the intervening years.

[11]This classification was based on products' placement into IQVIA's Uniform System of Classification pharmaceutical grouping codes

increase in non-circulatory products by 5% for those whose ER visits were after the January 1 cut-point. These results provide some confidence that myocardial infarction patients are prioritizing their circulatory medications regardless of whether their costs accrue to the outgoing plan-year or the upcoming plan-year. Meanwhile, they take advantage of lower prices for products unrelated to their circulatory system when their ER visit pushes them into lower cost phases of coverage.

The utilization impacts of the plan-year reset may also be different between branded and generic medications. Branded medications are more expensive and many have higher cost-sharing due to formulary design incentivizing generic utilization. Figure 3.8 shows the regression discontinuity data for branded and generic products in terms of total pharmaceutical expenditures in the 365 days following a myocardial infarction ER visit. No effect is found for generic medications, while an increase of $187 is found for total branded expenditures (95 % CI: $36 to $337). This corresponds to an increase in branded expenditures of 8% due to the cost accruals of the myocardial infarction ER visit.

3.5.2. Monthly Effects On Average Coinsurance Rate. Figure 3.9 shows the regression discontinuity plots for the average coinsurance rate that patients pay for pharmaceuticals for each thirty-day increment following an ER visit related to a myocardial infarction. The eligible data in the thirty-day regressions is restricted in two dimensions: Firstly, my extract contains data from July 8$^{\text{th}}$ of the previous plan-year to June 30$^{\text{th}}$ of the main plan-year. Secondly, for the regression discontinuity design to be consistently estimated using points far from the cut-point, the dependent variable must vary smoothly with the running variable: days from January 1$^{\text{st}}$. This is generally expected to be true unless the thirty-day increment for an injury date crosses into the following plan-year, in which case there is a discontinuity in the expected coinsurance rate. This limits the injury dates included in the analysis especially for thirty-day increments either close to day zero or to day 365.

Figure 3.10 graphically depicts the regression discontinuity coefficients estimated using both parametric and nonparametric approaches . These results show that entering the ER on January 1st leads to large decreases in average coinsurance rates relative to entering on December 31st. While the effect on average pharmaceutical coinsurance rate is largest around 2 to 5 months after the ER visit, the effect is statistically significant and substantial in the measurable timeframe from 1 to 11 months after the ER visit.

3.5.3. Parsing the Effects of Deductible and OOP Limit. I hypothesize that two factors lead to a decrease in average coinsurance rate: (1) the reduced costs related to the injuries transition patients from the deductible phase to the coinsurance phase, and (2) the reduced costs related to the injuries transition patients from the coinsurance phase to the OOP limit phase. To test these hypotheses, I evaluate whether there is a discontinuity in full-cost prescriptions at the cut-point to test the deductible-to-coinsurance phase transition and whether there is a discontinuity in zero-dollar prescriptions at the cut-point to test the coinsurance-to-OOP limit transition. Figure 3.11 shows the regression discontinuity coefficients of number of full-cost scripts relative to injury date for myocardial infarction patients. At the cut-point, there is a decrease in the number of full-cost prescriptions of about 0.5 prescriptions per thirty days in the days 30 to 119 following the ER visit (see Figure 3.12). This difference decreases down to 0.2 after 240 days after the ER visit. I hypothesize this is due to the December ER visit patients naturally passing through their deductible by this time.

Figure 3.13 shows the same plot for zero-dollar scripts relative to injury date for myocardial infarction patients. At the cut-point, there is an increase in the number of zero-dollar OOP prescriptions of about 1 prescription per 30 days following the injury (see Figure 3.14). This decreases throughout the year following the ER visit but not as quickly as the full-cost claims do. Even as late as 300 days after the ER visit, the January ER visits still have 0.75 more zero-dollar claims per thirty days than the December ER visits do. This slower decay

may be due to the low likelihood of naturally reaching the OOP limit for the December ER visit patients.

To answer the question as to which kink in the insurance contract[12] is driving the decrease in average coinsurance rate, it is important to remember that full-cost claims are much farther from the average coinsurance rate than are zero-dollar claims. Since the average coinsurance rates in the first few months after the ER visit for December ER visits are about 33%, and the zero-dollar increase (about 1 additional zero-dollar claim per 30 days) is about twice as big as the full-cost decrease (about 0.5 fewer full-cost claims per 30 days) then the the effects on the average coinsurance rate are approximately equal. Theoretically, the impact of the deductible phase kink and the OOP limit phase kink depends on the severity of the condition precipitating the ER visit. In this case, since myocardial infarctions are generally severe enough to move some patients all the way to the OOP limit, that kink is more impactful than it would be for other less severe injuries such as broken bones and sprains. In those cases, the deductible kink would likely dominate the cost impact.

3.5.4. Monthly Pharmacy Expenditures and the Extensive Margin. When examining the total expenditures at the 30-day intervals, there is more variance in the data due to a smaller sample of claims in a shorter time period which offsets the larger magnitude in average coinsurance rates. The data is shown in Figure 3.15 and the global parametric and local nonparametric coefficients are shown in Figure 3.16. Most data points are near zero although there are two outlier 30-day periods which show a marginally significant increase in pharmaceutical expenditures in both regression approaches.

To understand whether there are any effects of the plan-year reset on the extensive margin, in Figure 3.17, I examine the proportion of patients who did not fill any pharmaceutical claims after their ER visit. There are three reasons why people may not be observed to fill any medication in the data: (1) they use a pharmacy not captured in the IQVIA data

[12]the kink between the deductible and coinsurance phase, or the kink between the coinsurance phase and OOP limit

sample, (2) they are not adherent to their medication or are not prescribed any, (3) they may have died. The first observation to make is that there is no measured or visible impact of the plan-year reset in the extensive margin for filling pharmaceutical medications after myocardial infarction ER visits. Next, the results for kidney stone-related ER visits are included to provide suggestive evidence as to the magnitude of these reasons. It is plausible that kidney stone patients go to similar pharmacies as myocardial infarction patients and both patient groups are very likely to be prescribed medications after their ER visit. It is difficult to know about how adherent the two patient groups are to the medications prescribed to them. However, it is much more likely that myocardial infarction patients die due to their condition than kidney stone patients to their condition. The difference in the proportion of patients without any pharmaceutical claims between the two groups is indirect evidence about the proportion of myocardial patients who die shortly after their ER visit. That difference is measured to be 4.7%. This is consistent with estimates of myocardial infarction mortality in the medical literature which finds that the in-hospital myocardial infarction mortality rate is 4.6% (McNamara et al. 2016). Due to the insignificant discontinuity in this data and no direct mortality measurements, this study can provide no evidence to suggest that the timing of a myocardial infarction ER visit relative to the January 1st cut-point affects mortality.

3.5.5. Billing Delay. As was mentioned in Section 3.4, there is often a large delay between the date of a myocardial infarction ER visit and the date the patient is sent the bill (see Figure 3.3). One may wonder how these delays affect the pricing dynamics found in this study. If the bill has not been calculated and negotiated by the hospital and the insurer, then the insurer may be unable to transition patients into their correct phase of coverage. For patients who have not yet been sent their ER bill, there is no statistically significant evidence of lower average coinsurance rates, however the data cannot reject the idea that it may be possible for some insurers to adjust OOP costs before the bill is sent to the patient. Figure 3.18 shows the effect of the discontinuity on the average pharmaceutical coinsurance

rates for patients who have been sent their bills. When restricting the analysis to exclude patients who have not been sent their bill by the beginning of each 30-day period, the effect on average coinsurance is slightly larger in the first two months following the ER visit, but the effect of this exclusion is small relative to the statistical error in the analysis. The point estimate for the 30-59 day period post-MI is 0.8% lower than in the full sample. The effect is negligible for time periods after 120 days because most patients will have received their bill by 120 days after their ER visit.

Delays in billing may reduce the effect size of the plan-year reset on pharmacy OOP costs. This is to the detriment of patients who have their myocardial infarctions early in the plan-year and have a long delay in their billing unless the excess pharmacy OOP costs that they are exposed to are credited to the patient after the ER bill is settled. I was unable to find any anecdotal data of insurers reimbursing patients for excess OOP costs between a medical event and the finalized bill. Whether or not they are reimbursed for their excess costs over their contract due to the delayed ER bill, they are likely still making pharmacy purchasing decisions based on the higher prices they are shown at the pharmacy. In the context of my study, these patients with severely delayed bills that prevent the average coinsurance rate from changing can be viewed as non-compliers.

3.6. Discussion

This analysis found large effects of the plan-year reset on average coinsurance rates for pharmaceuticals for privately-insured patients after visiting the ER with a myocardial infarction diagnosis regardless of estimation methods. These effects constitute an ex-post inequity for patients who have their myocardial infarction near the end of the calendar year compared to those who have their myocardial infarction in the beginning of the year. Using a global parametric regression, the effects on total pharmaceutical expenditure and other pharmacy utilization metrics related to filled prescriptions are statistically significant and similar to those found for outpatient expenditures in the RAND Health Experiment. With

a patient population as severely ill as those with a recent myocardial infarction, it is not unexpected to find very low elasticities of demand especially for high-value generic products that treat that specific life-threatening condition.

While OOP limits affect a small proportion of all privately-insured consumers, there are many sub-populations for whom it is impactful and increases the non-linearity of their prices. Patients who are admitted to the ER with myocardial infarctions are one such sub-population and this study finds that the OOP limit is similarly impactful as the deductible. This means that using fuzzy regression discontinuity designs with a first-stage related to passing through the deductible phase may be significantly biased upwards because passing through the deductible is positively correlated to reaching the OOP limit.

A key objective of this line of research is to develop a method to precisely measure patient price sensitivity for many different patient populations and with many different healthcare products. The pharmacy and hospital data used in this study provide a wide sample of U.S. ER visits and pharmacy claims that are likely more representative of the overall population than the data from a subset of insurers or employers used by other studies. However, the limitations of this data to observe important patient characteristics such as deductible and OOP limit thresholds, as well as complete in-network healthcare spend prevent me from being able to estimate a counterfactual price based on the cost of the ER visit. The ability to do this would allow researchers to differentiate which patients are likely compliers (those with high deductibles, high ER costs, and low OOP limits) from likely non-compliers who have more linear insurance contracts. This approach would be similar to the one used by (Chandra et al. 2024) which they found to provide much needed additional power.

Future studies should endeavor to increase this power in other ways as well. The simplest approach is to gather more ER visits. The desire to avoid regression contamination due to the COVID pandemic in 2020 and 2021 reduced this study's sample size by a third, and every year, more data is being collected that could be aggregated to reduce variance.

Another approach would be to examine outcomes outside of pharmaceuticals such as hospital admissions, outpatient visits, and mortality that might have different effect sizes. Finally, aggregating other ER visit reasons could provide additional power at the cost of specificity.

As long as the substantial non-linearity of health insurance contracts persist in the U.S. healthcare system, this "as-if-random" variation in prices based on the accrual of ER visit costs will continue to be a useful natural experiment to understand how prices affect health outcomes for those with unexpected healthcare shocks near the plan-year reset. It is also possible to extend this analysis to Medicare where deductibles exist but are much smaller and where OOP limits are being instituted for pharmaceuticals in the Inflation Reduction Act of 2022.

Finally, to answer the title question, depending on a person's insurance contract, it is better to have a heart attack early in the plan-year so that the related ER and hospitalization costs accrued push them through their deductible and coinsurance phases of coverage to achieve lower healthcare prices in the cheaper phases of coverage for the remainder of the plan-year. Alternatively, it may be preferable for a patient who has already reached their out-of-pocket limit to have a heart attack right after they reached that limit so that the whole near-term cost of the ER and hospitalization experience is borne by the insurer. Any strongly non-linear insurance contract will have these ex-post inequities, but as insurance contracts become more non-linear due to policy-driven incentives for large deductibles and regulations that drive binding out-of-pocket limits, policymakers should be cautious of the magnitude of these inequities in the U.S. insurance and healthcare system.

3.7. Tables

TABLE 3.1. ER Diagnosis and Procedure Codes for Myocardial Infarctions with Private Insurance

Characteristic	**Overall**, N = 57,031[†]	**Left Side**, N = 26,940[†]	**Right Side**, N = 30,091[†]
ER Diagnosis Description			
Non-ST elevation (NSTEMI) myocardial infarction	28,343 (50%)	13,248 (49%)	15,095 (50%)
Secondary Diagnosis: myocardial infarction	13,891 (24%)	6,576 (24%)	7,315 (24%)
ST elevation (STEMI) myocardial infarction of unspecified site	11,363 (20%)	5,451 (20%)	5,912 (20%)
ST elevation (STEMI) myocardial infarction involving other coronary artery of inferior wall	2,385 (4.2%)	1,205 (4.5%)	1,180 (3.9%)
Acute myocardial infarction, unspecified	1,049 (1.8%)	460 (1.7%)	589 (2.0%)
ER Procedure Description			
Code 99285: Emergency department visit, high level decision making	53,436 (94%)	25,296 (94%)	28,140 (94%)
Code 99284: Emergency department visit, moderate level decision making	3,107 (5.4%)	1,426 (5.3%)	1,681 (5.6%)
Code 99283: Emergency department visit, low level decision making	488 (0.9%)	218 (0.8%)	270 (0.9%)

[†] n (%)

Notes: ER visits are classified as myocardial infarctions if they are attached with ICD10 diagnosis codes related to myocardial infarctions which includes all ICD10 codes beginning with I50. The distribution of primary diagnosis descriptions is shown; if the myocardial infarction diagnosis is in a secondary position, the visit is grouped into the "secondary diagnosis" category. Procedure codes are entered into the system associated with each ER visit that are related to the effort and complexity required for care. The data is separated into two groups: before the January 1 cut-point (left side) and after (right side). Source: IQVIA Medical Claims.

TABLE 3.2. Demographics for Myocardial Infarction ER Visits with Private Insurance

Characteristic	**Overall**, N = 57,031[1]	**Left Side**, N = 26,940[1]	**Right Side**, N = 30,091[1]
Age Group			
0-18	123 (0.2%)	54 (0.2%)	69 (0.2%)
19-44	5,956 (10%)	2,729 (10%)	3,227 (11%)
45-64	38,018 (67%)	17,907 (66%)	20,111 (67%)
65+	12,934 (23%)	6,250 (23%)	6,684 (22%)
Sex			
F	18,266 (32%)	8,745 (32%)	9,521 (32%)
M	38,765 (68%)	18,195 (68%)	20,570 (68%)

[1] n (%)

Notes: ER visits are classified as myocardial infarctions if they are attached with ICD10 diagnosis codes related to myocardial infarctions which includes all ICD10 codes beginning with I50. Ages are calculated based on patient birth year, and sex is recorded as sex defined at birth. Patients whose age is 65 or over are excluded from regressions to avoid Medicare contamination. Source: IQVIA Medical Claims.

TABLE 3.3. Regression Discontinuity Coefficient for Average Pharmaceutical Coinsurance Rate over 365-Days following Myocardial Infarction ER Visit

	BIC-Chosen Parametric		Local-Linear Nonparametric	
Estimated Beta	-0.054^{***}	-0.051^{***}	-0.037^{***}	-0.035^{***}
(Std. Err.)	(0.0048)	(0.0051)	(0.012)	(0.013)
Polynomial Order	1	1	1	1
Coefficients Included	Yes	No	Yes	No
Bandwidth	177, 180	177, 180	60, 63	56, 58
Kernel	Uniform	Uniform	Triangle	Triangle
Num. obs.	48,137	48,136	16,521	15,360

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$

Notes: Regression discontinuity coefficients from Equations 1 and 2 estimated with and without covariates. The parametric regressions include myocardial infarction ER visits from July 8 to June 30 surrounding the January 1 cut-point. The nonparametric regressions select separate MSE-optimal bandwidths for each side based on the same data and restrict estimation to those data with a triangle kernel. Cluster-robust standard errors are shown. Source: IQVIA Medical and Pharmacy Claims.

TABLE 3.4. Regression Discontinuity Coefficient for Total Pharmaceutical Expenditure over 365-Days following Myocardial Infarction ER Visit

	BIC-Chosen Parametric		Local-Linear Nonparametric	
Estimated Beta	\$198**	\$198*	\$158	\$45
(Std. Err.)	(88)	(108)	(245)	(287)
Polynomial Order	1	1	1	1
Coefficients Included	Yes	No	Yes	No
Bandwidth	177, 180	177, 180	52, 41	51, 60
Kernel	Uniform	Uniform	Triangle	Triangle
Num. obs.	56,665	56,687	15,025	17,637

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$

Notes: Regression discontinuity coefficients from Equations 1 and 2 estimated with and without covariates. The parametric regressions include myocardial infarction ER visits from July 8 to June 30 surrounding the January 1 cut-point. The nonparametric regressions select separate MSE-optimal bandwidths for each side based on the same data and restrict estimation to those data with a triangle kernel. Cluster-robust standard errors are shown. Source: IQVIA Medical and Pharmacy Claims.

3.8. Figures

FIGURE 3.1. Density of Visits to the ER with a Myocardial Infarction Diagnosis

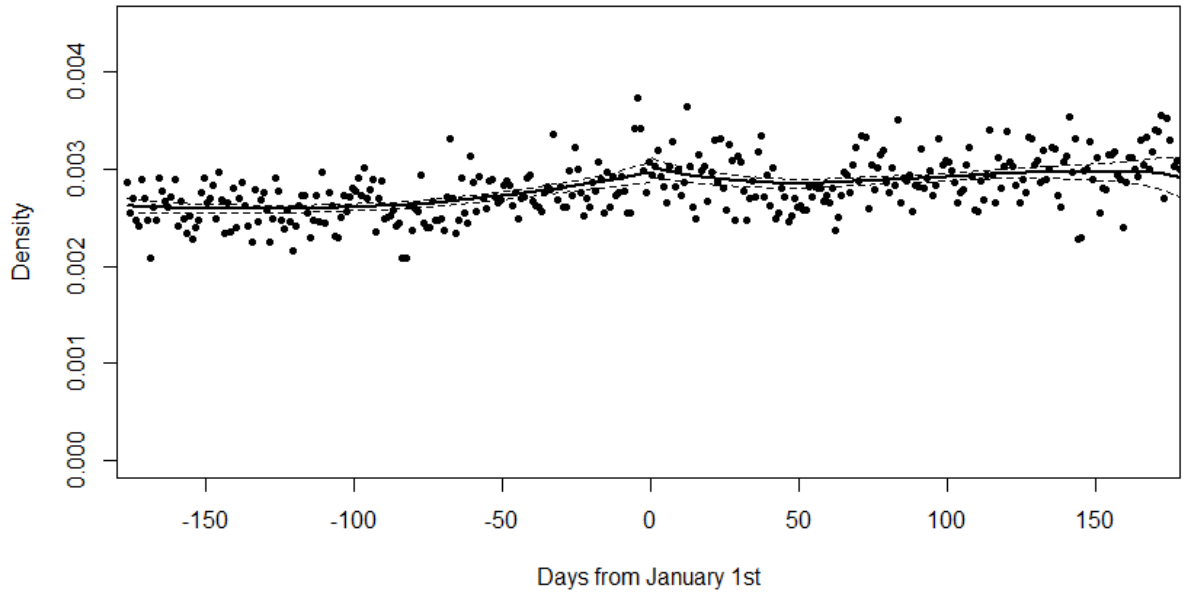

Notes: Visits are included in the analysis if they are paid for by a private insurer and are attached to a myocardial infarction diagnosis. The solid line is the predicted density and the dashed lines represent the 95% confidence interval. Source: IQVIA Medical Claims.

FIGURE 3.2. Example Insurance Contracts with 2022 Maximum OOP Limits

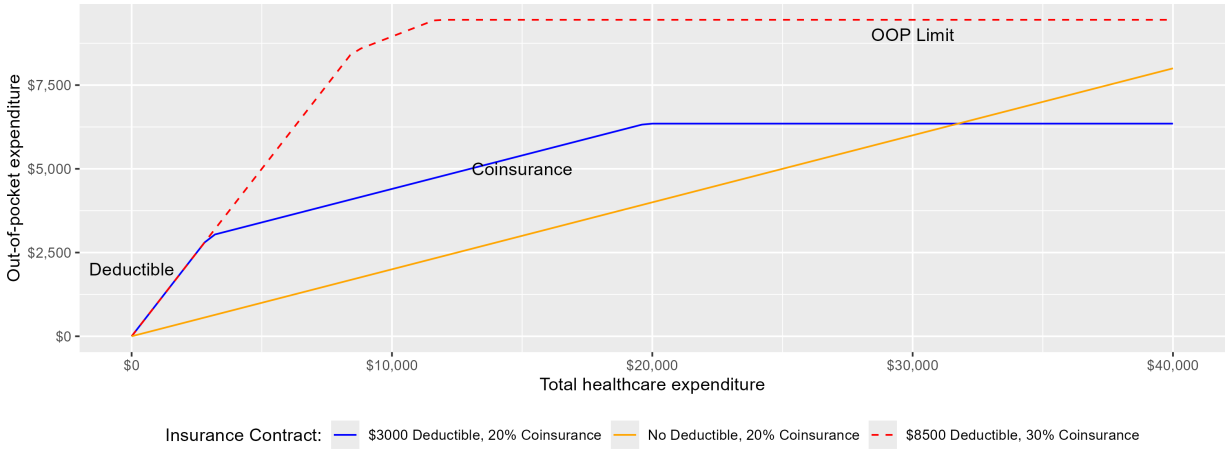

Notes: An illustrative example of the relationship between total healthcare expenditure and out-of-pocket expenditure paid by the patient fpr three hypothetical plans: (Blue) a high-deductible healthcare plan with a $3,000 deductible and an average 20% coinsurance during the coinsurance/copay coverage phase with an OOP limit of $6,100. (Red Dashed) a very high-deductible healthcare plan with a $8,500 deductible and an average 30% coinsurance during the coinsurance/copay coverage phase with an OOP limit equal to the 2022 maximum OOP limit of $9,450. This may be similar coverage to a Bronze level exchange plan. (Orange) a healthcare plan with no deductible and 20% average coinsurance during the coinsurance phase without a visible OOP limit.

FIGURE 3.3. The Wait Time Between Myocardial Infarction ER Visit and Date of Corresponding Bill

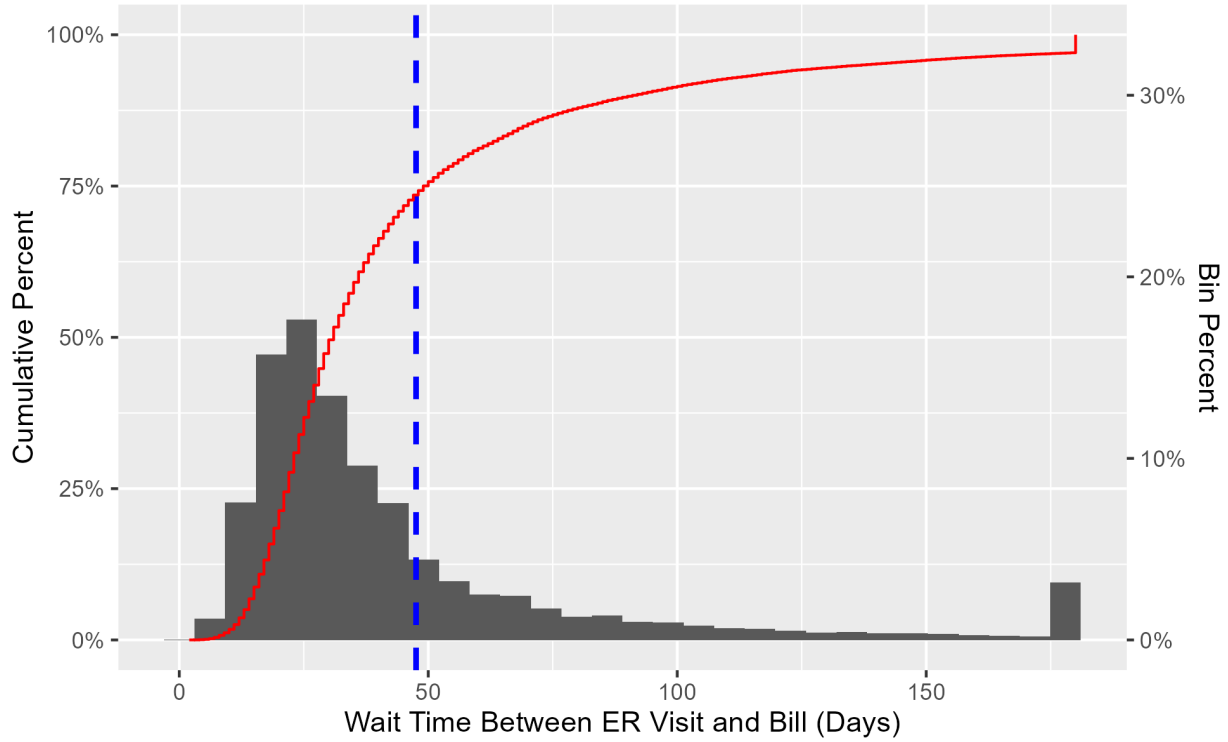

Notes: Wait times are censored at 180 days with 3% of ER visits censored. The billing wait time is calculated based on the ER visit procedure and any subsequent care may have a different billing delay.
Source: IQVIA Medical Claims

FIGURE 3.4. The Mean Hospital / Clinic Out-of-Pocket Costs for Care Relative to the Myocardial Infarction ER Visit Date

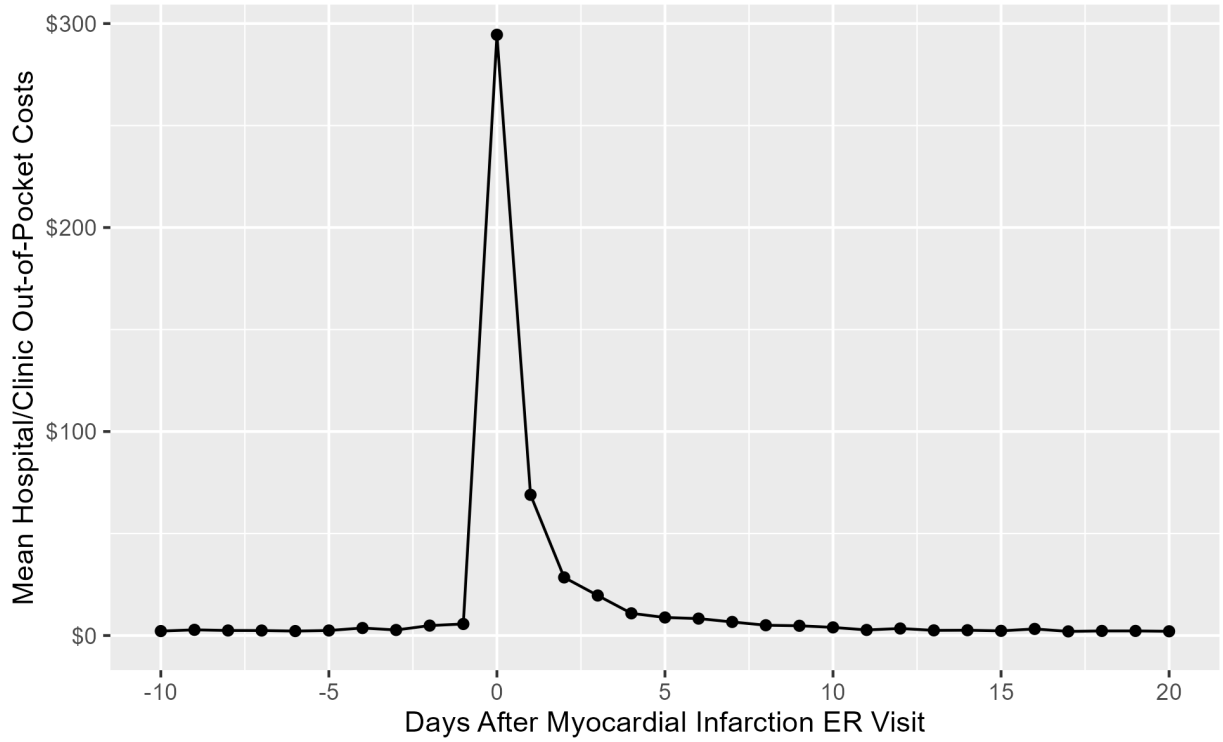

Notes: All hospital and clinic out-of-pocket costs are aggregated for each patient in the study. The mean of these patients out-of-pocket costs is then calculated relative to the ER date. Source: IQVIA Medical Claims

FIGURE 3.5. Relationship between Non-Outcome Variables and ER Date: Age, Sex, and Prescription Volume and Prescription Expenditures in the 90 Days Preceding the ER Visit

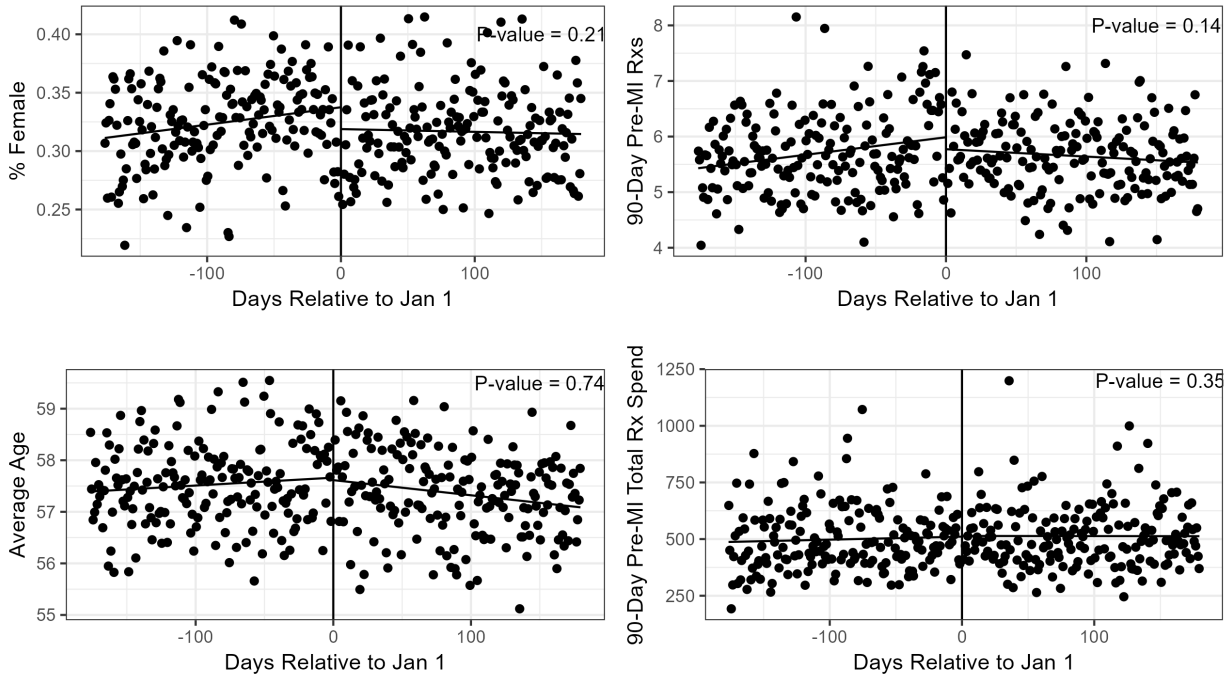

Notes: Scatter plots of non-outcome variables: sex, age, and, prescriptions filled and total prescription expenditures in the 90-day period before the myocardial infarction ER visit and age relative to the January 1 cut-point. Global linear-interaction parametric regressions are shown and p-values are calculated. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.6. Full-Year Coinsurance and Total Pharmaceutical Expenditures

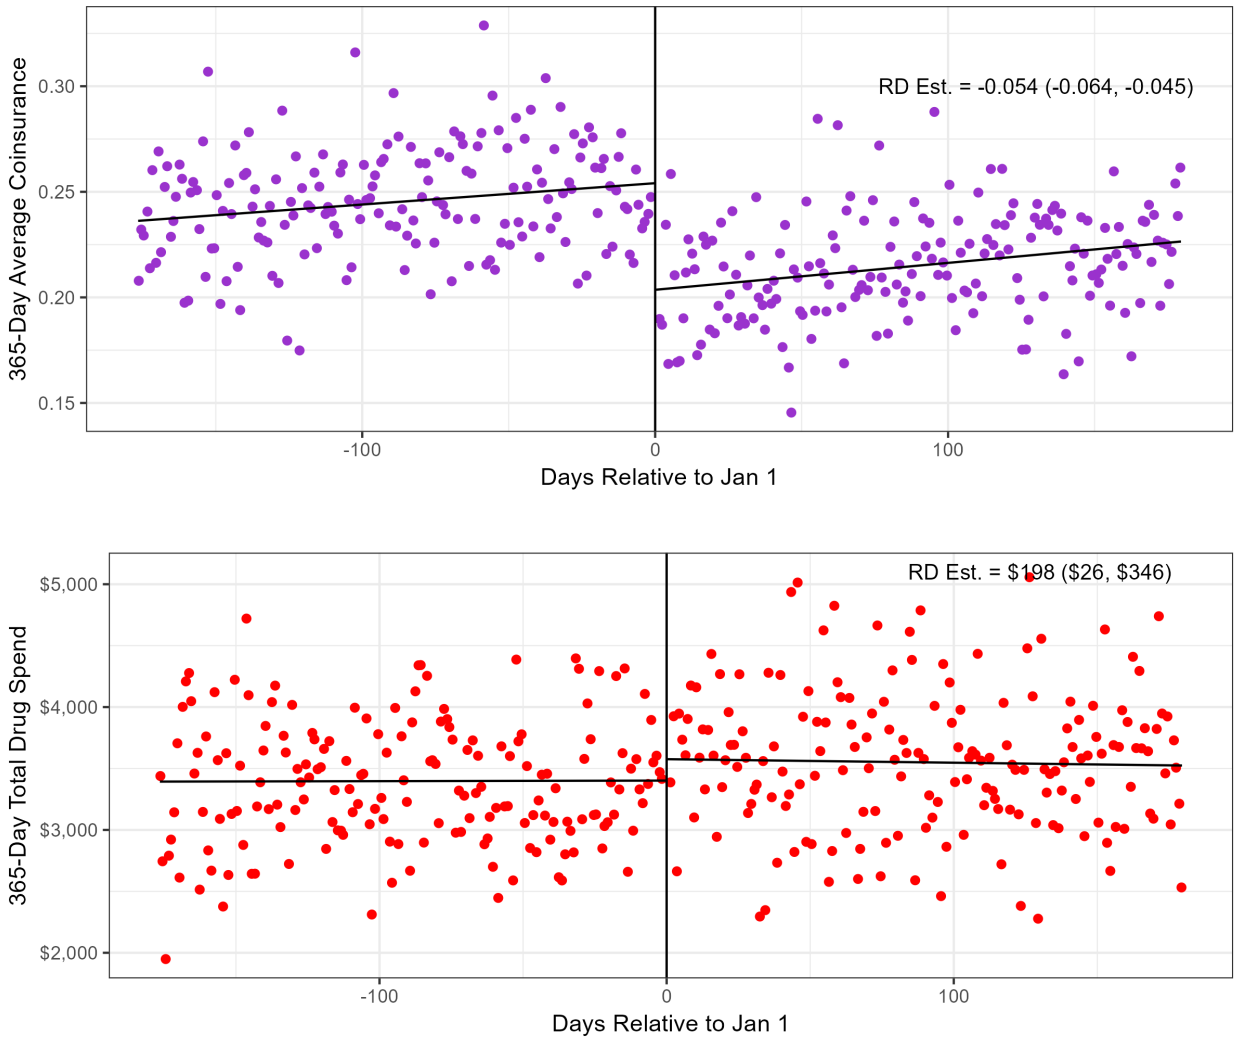

Notes: (Top) Dependent variable is the 365-day post ER visit average coinsurance rate = (Total Pharmaceutical OOP Spending) /(Total Pharmaceutical Expenditures). (Bottom) Dependent variable is the 365-day post ER visit total pharmaceutical expenditures. Global parametric regression with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.7. Full-Year Number of Rxs Filled

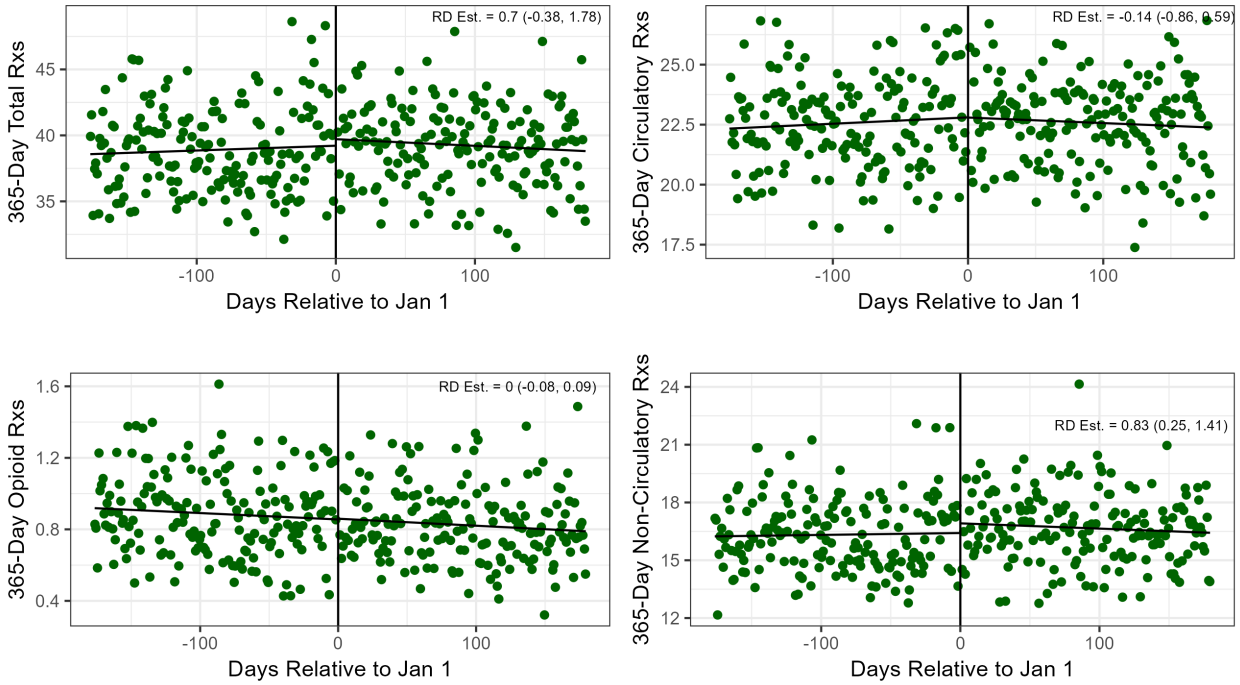

Notes: (top-left) All prescriptions. (top-right) Prescriptions related to circulatory conditions based on Uniform System of Classification (USC) Codes. (bottom-left) Prescriptions for opioid products. (bottom-right) Prescriptions for all products not related to circulatory conditions based on USC Codes. All variables are calculated for 365-days post ER visit and the global parametric regressions with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.8. Full-Year Branded and Generic Total Pharmaceutical Expenditures

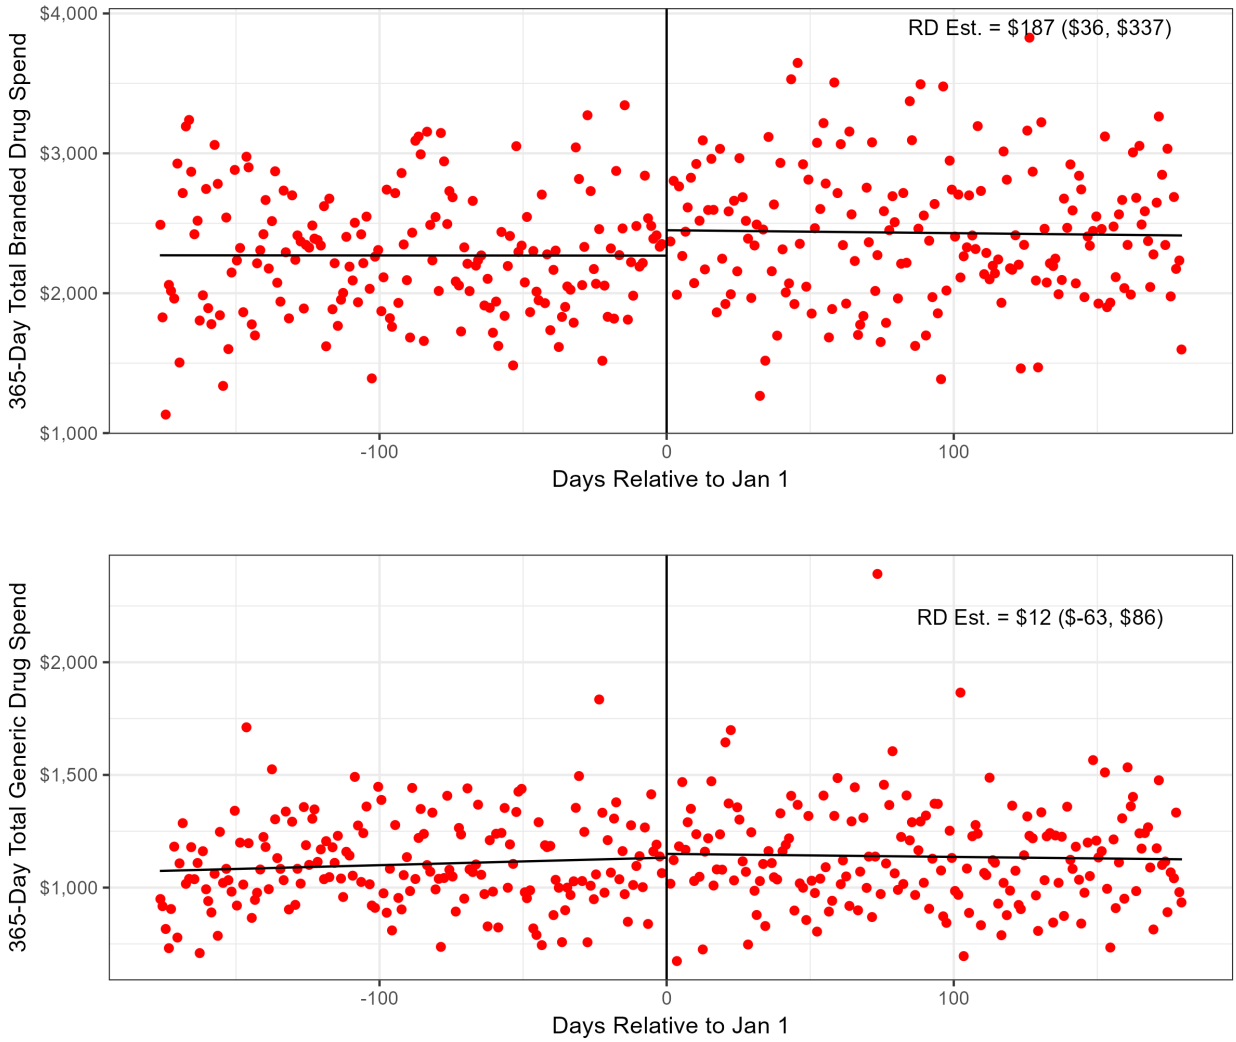

Notes: (Top) Dependent variable is the 365-day post ER visit total pharmaceutical expenditures on branded products. (Bottom) Dependent variable is the 365-day post ER visit total pharmaceutical expenditures on generic products. Global parametric regression with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.9. Myocardial Infarction Average Coinsurance Rate RD Plots

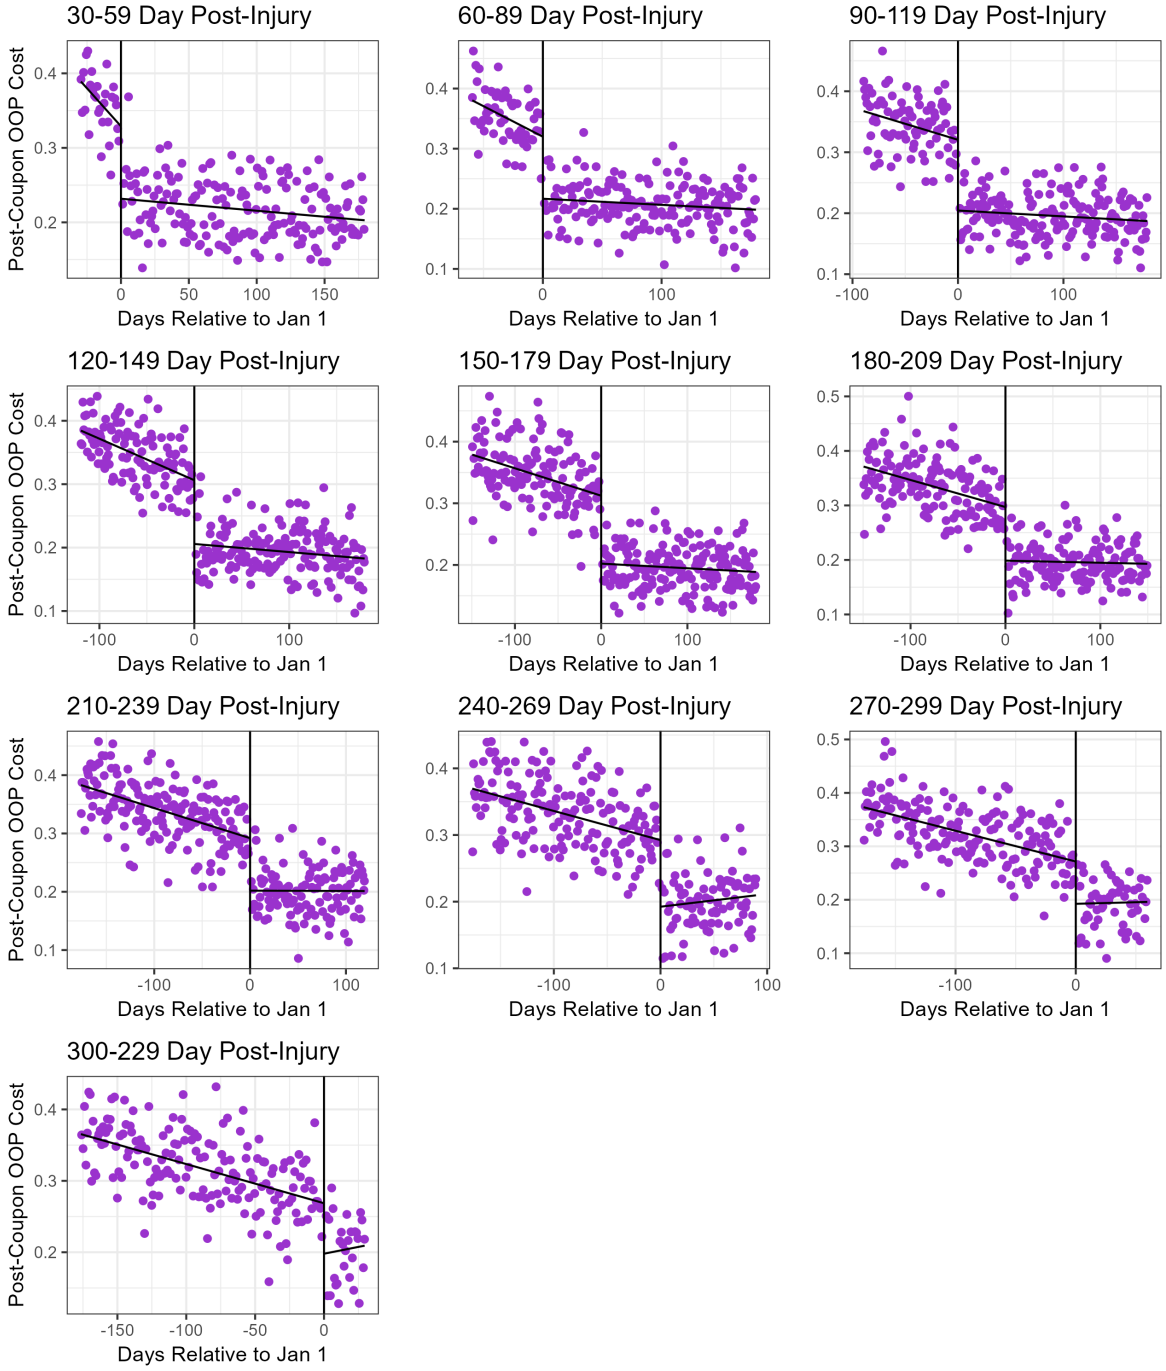

Notes: Scatter plots of the average coinsurance rate (total amount paid by patient / total reimbursement) relative to the ER visit date. Global parametric regression with BIC-selected interaction terms are shown. To observe the impact of an ER visit on coinsurance rate over the course of the following year, each graph represents the average coinsurance rate for the claims a certain length of time after the ER visit date. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.10. Myocardial Infarction Average Coinsurance Rate RD Coefficients

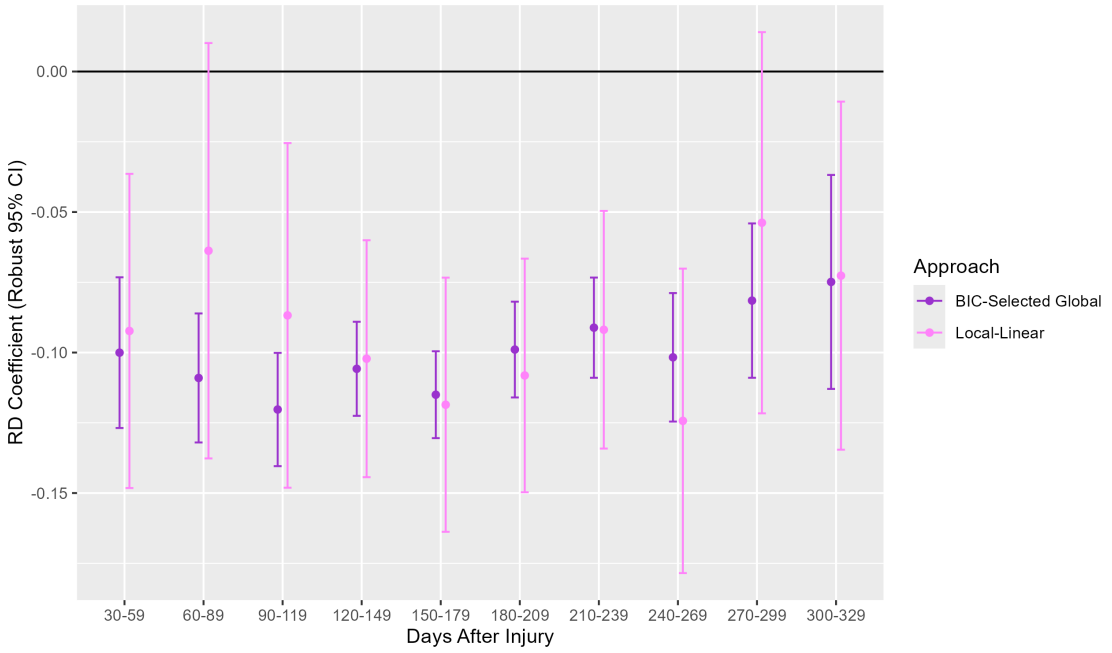

Notes: Coefficients from regression discontinuity estimates of Equations 1 and 2. The vertical axis represents the estimated difference at the cut-point in average coinsurance rate for each thirty-day period. The local linear regression uses optimal bandwidths and triangle kernel while the global parametric regressions are the same as shown in Figure 3.9. 95% robust confidence intervals are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.11. Myocardial Infarction Full-Cost Claims RD Plots

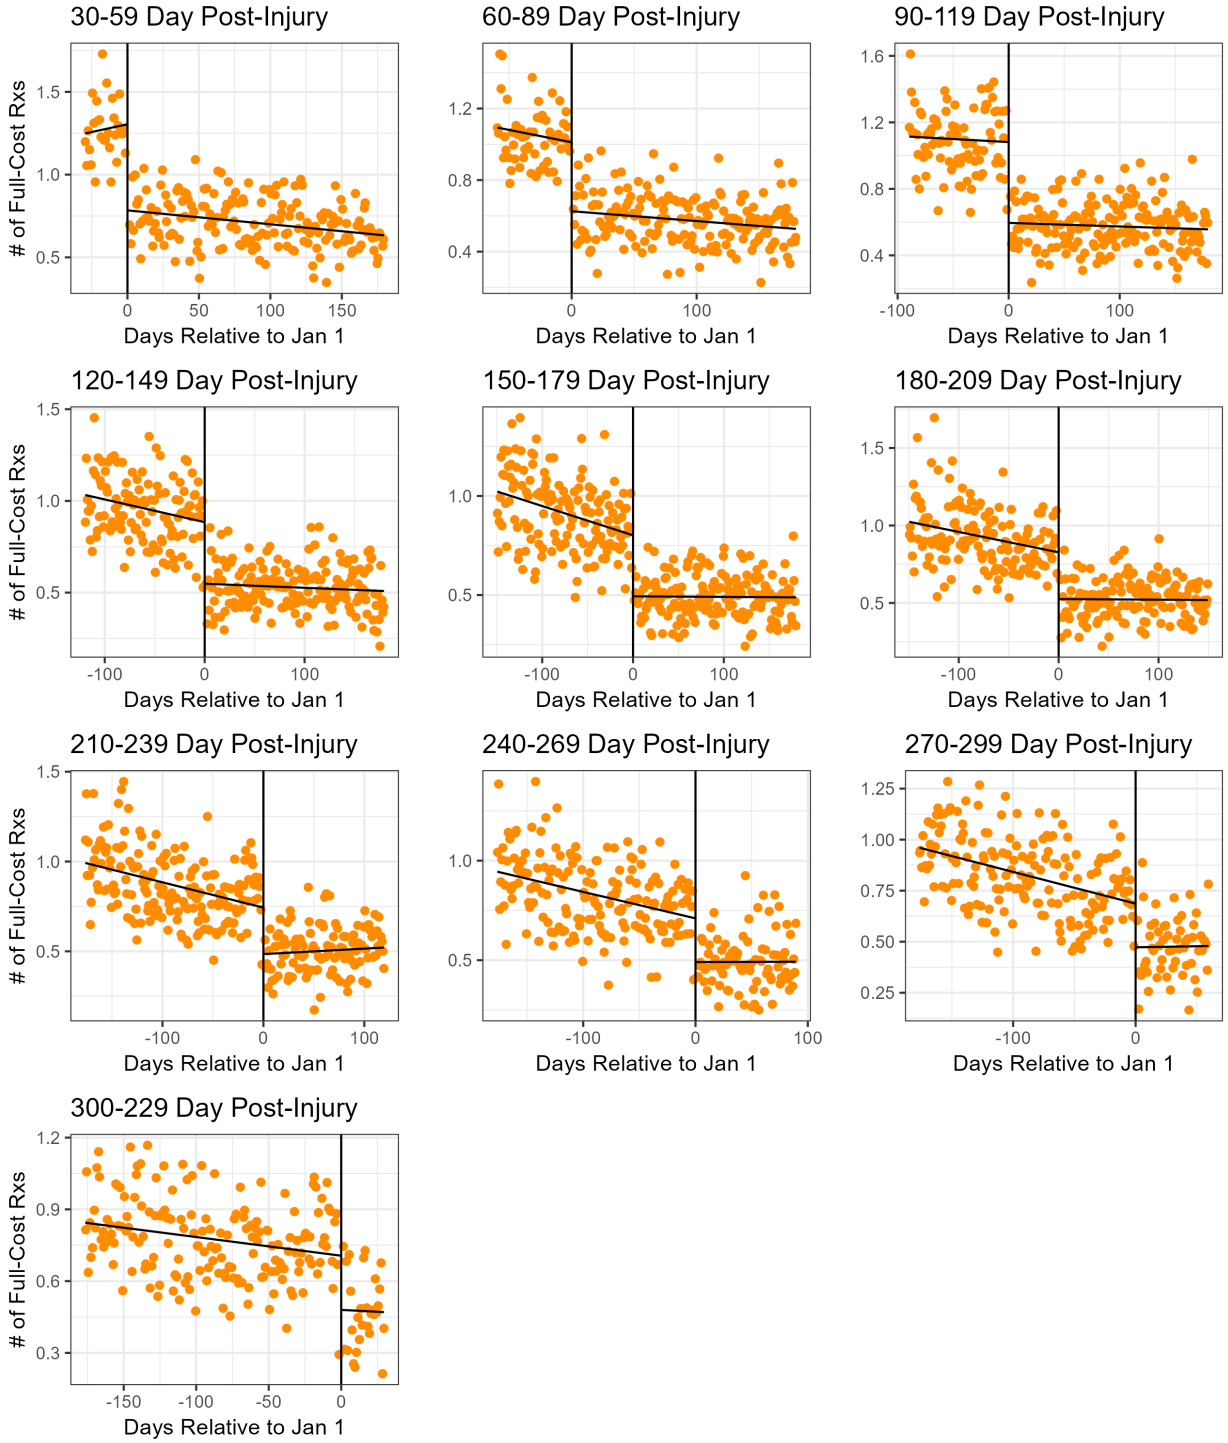

Notes: Scatter plots of the number of full-cost claims for each thirty-day period relative to the ER visit date. Global parametric regression with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.12. Myocardial Infarction Full-Cost Claims RD Coefficients

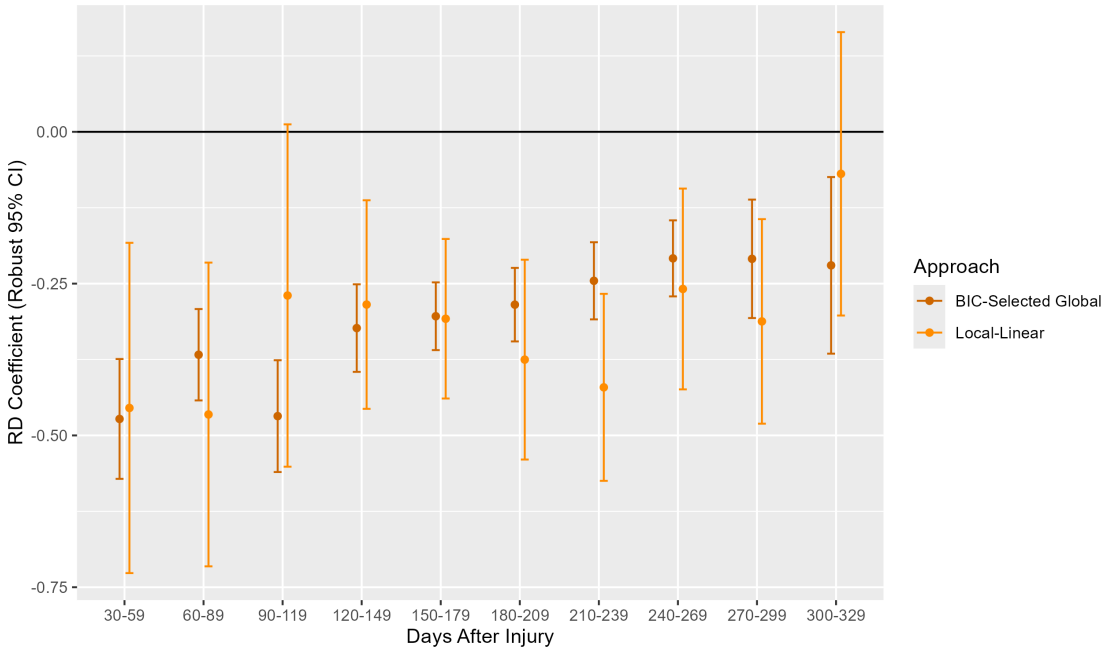

Notes: Coefficients from regression discontinuity estimates of Equations 1 and 2. The vertical axis represents the estimated difference at the cut-point in full-cost claims (claims where the patient pays out-of-pocket the full price of the product) for each thirty-day period. These full-cost claims among privately insured consumers are often related to the deductible phase of coverage. The local linear regression uses optimal bandwidths and triangle kernel while the global parametric regressions are the same as shown in Figure 3.11. 95% robust confidence intervals are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.13. Myocardial Infarction Zero-Dollar Claims RD Plots

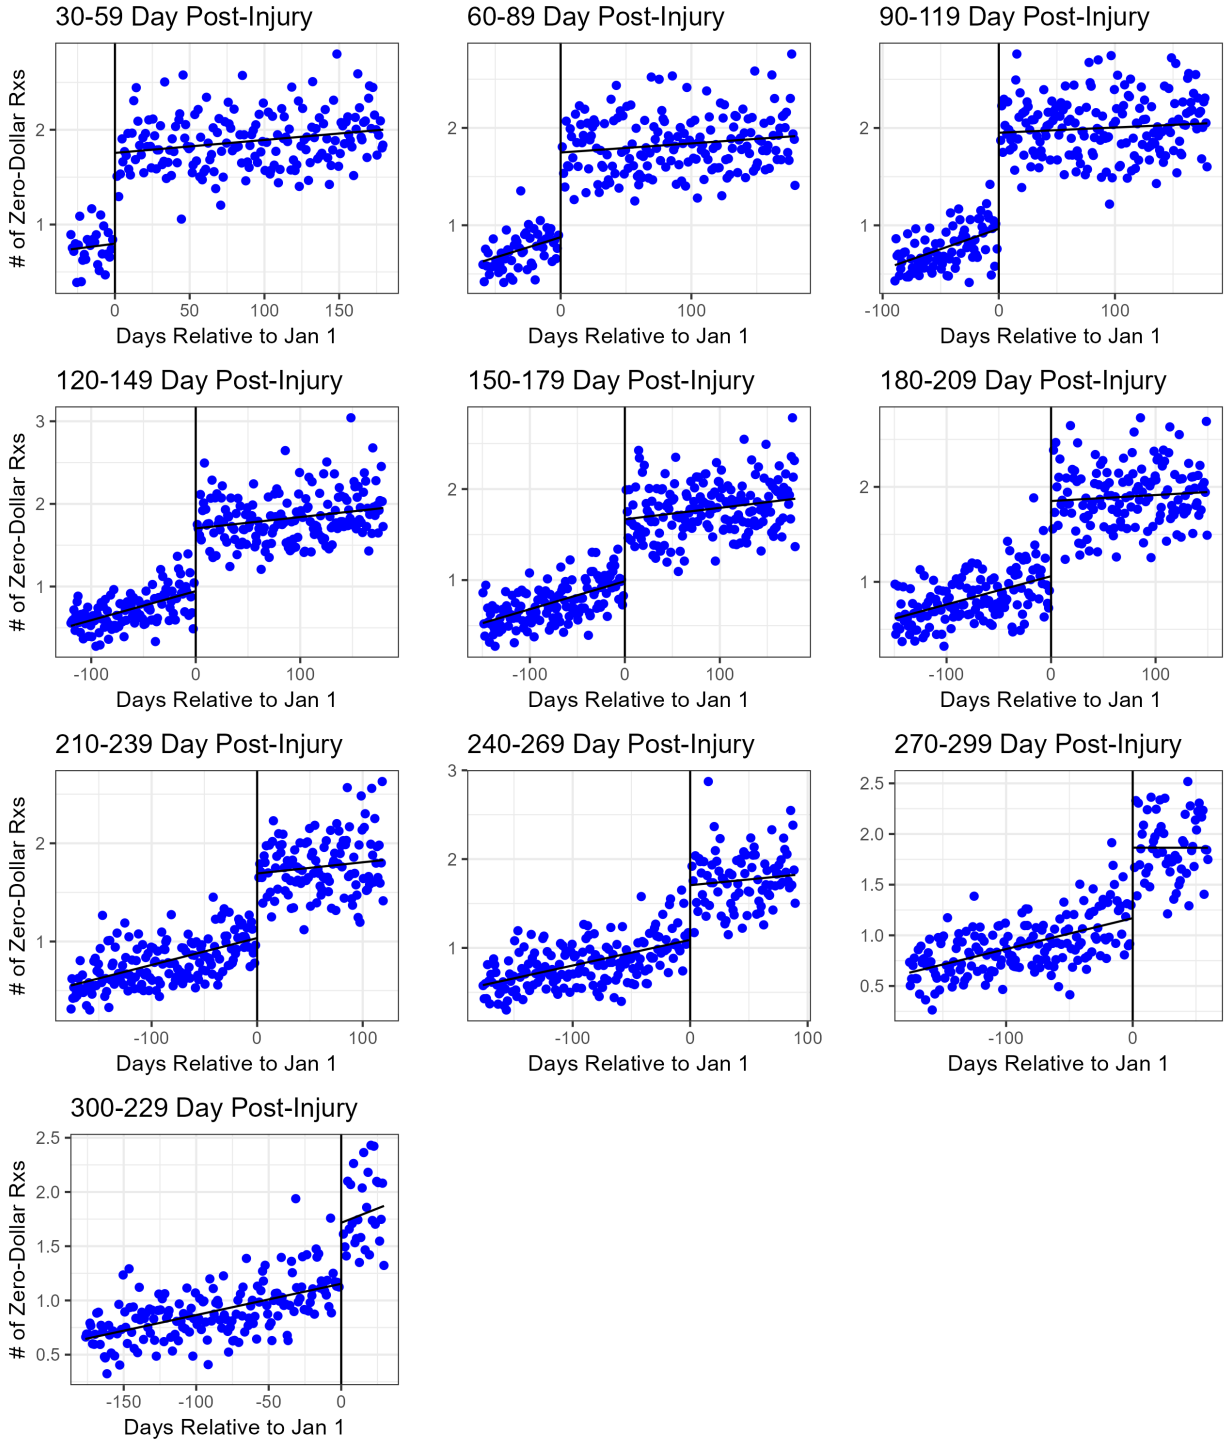

Notes: Scatter plots of the number of zero-cost claims for each thirty-day period relative to the ER visit date. Global parametric regression with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.14. Myocardial Infarction Zero-Dollar Claims RD Coefficients

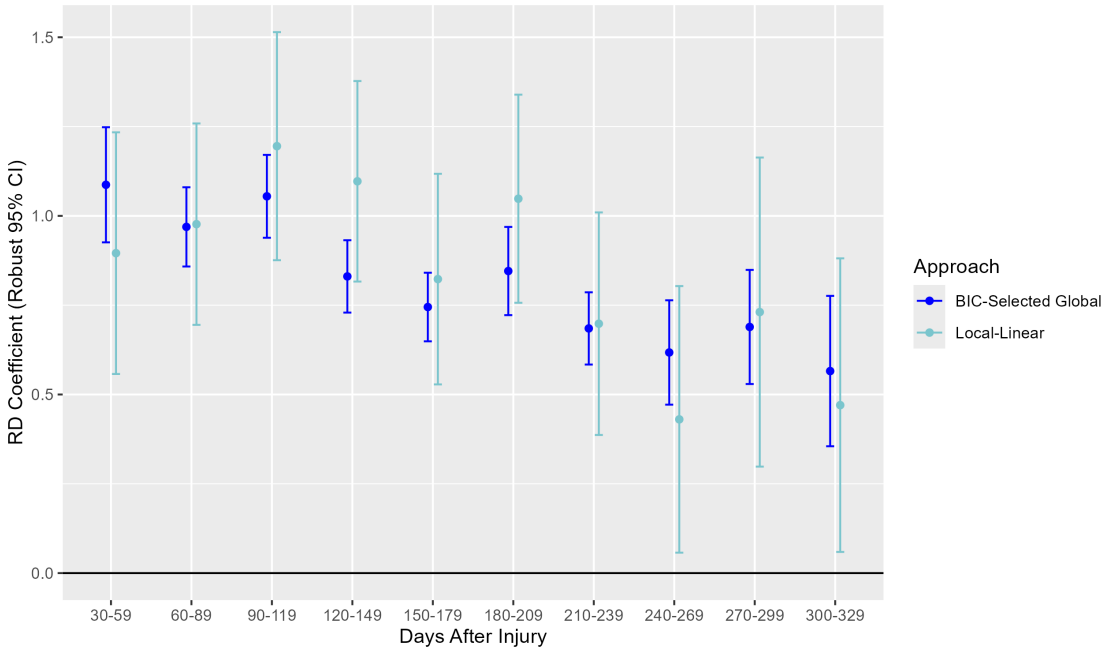

Notes: Coefficients from regression discontinuity estimates of Equations 1 and 2. The vertical axis represents the estimated difference at the cut-point in zero-cost claims (claims where the patient pays nothing out-of-pocket for the product) for each thirty-day period. These zero-cost claims among privately insured consumers are often related to the OOP limit phase of coverage. The local linear regression uses optimal bandwidths and triangle kernel while the global parametric regressions are the same as shown in Figure 3.13. 95% robust confidence intervals are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.15. Myocardial Infarction Average Total Expenditures RD Plots

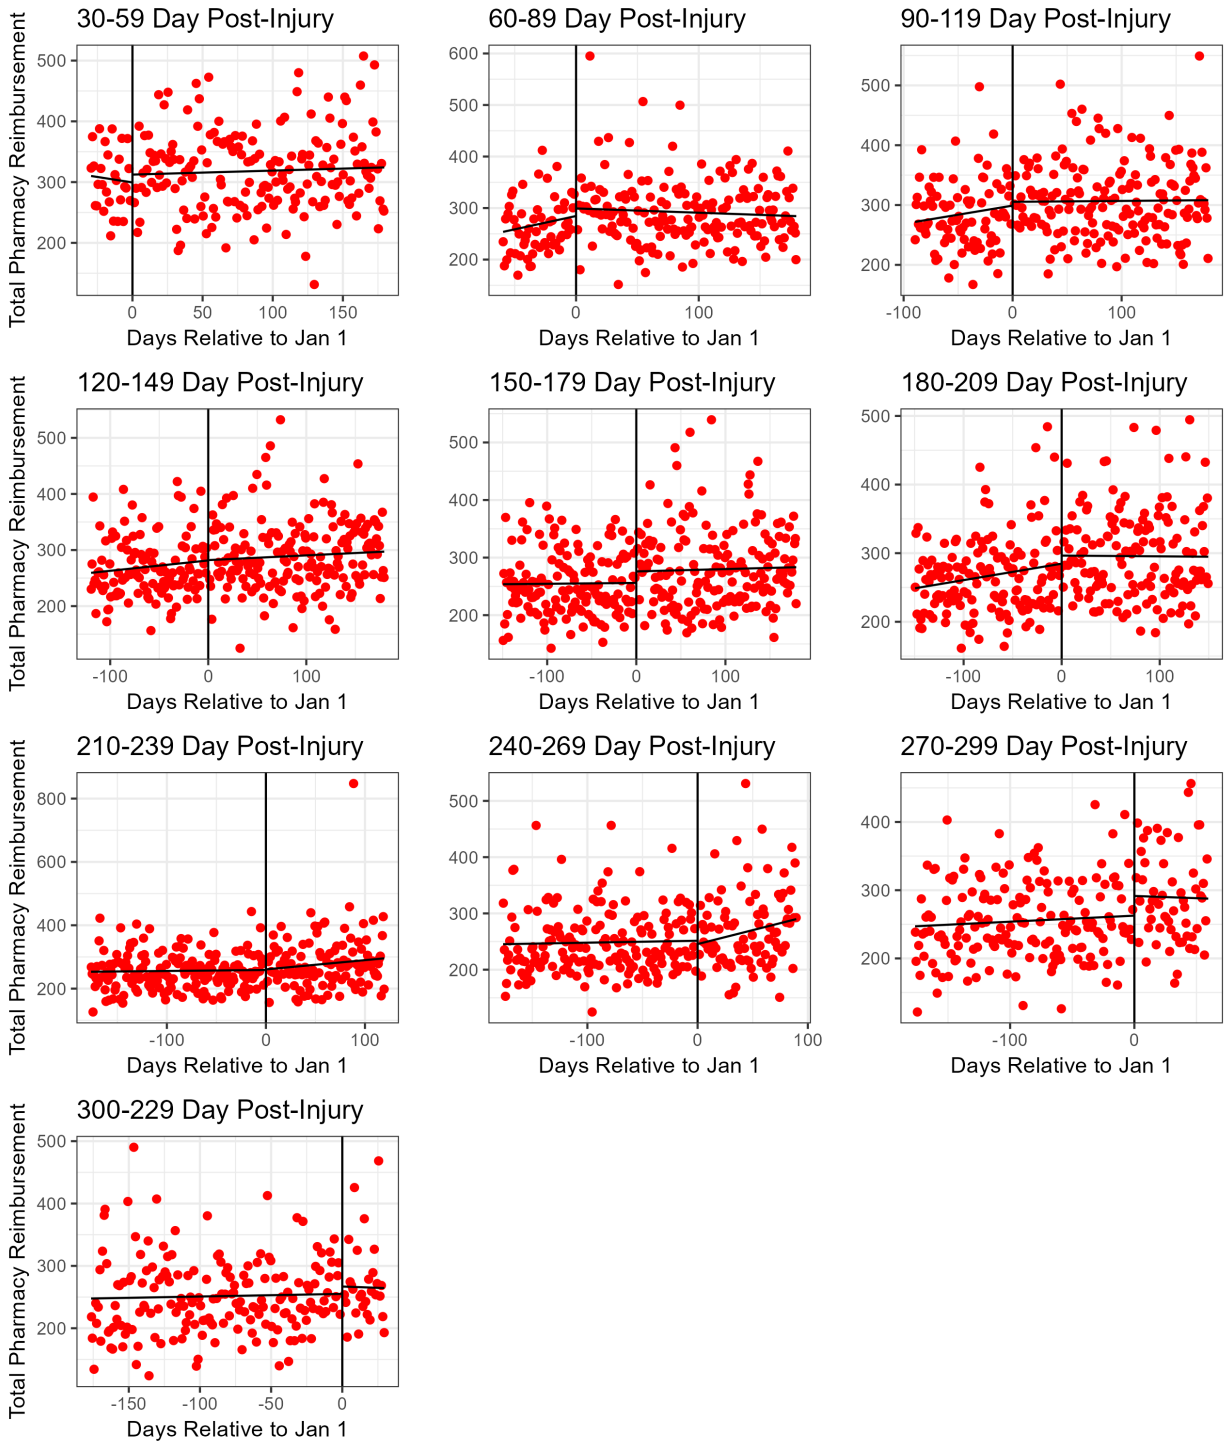

Notes: Scatter plots of the average total reimbursement (total expenditures including patient and insurer payments) relative to the ER visit date. Global parametric regression with BIC-selected interaction terms are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.16. Myocardial Infarction Average Total Expenditures RD Coefficients

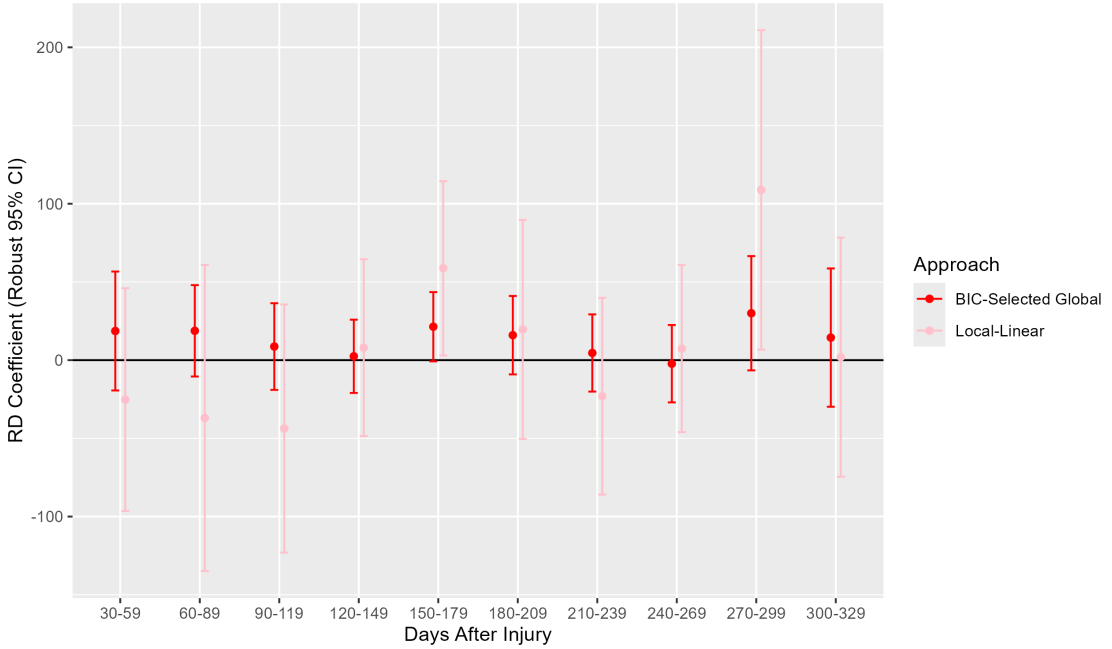

Notes: Coefficients from regression discontinuity estimates of Equations 1 and 2. The vertical axis represents the estimated difference at the cut-point in the average total reimbursement (total expenditures including patient and insurer payments) for each thirty-day period. The local linear regression uses optimal bandwidths and triangle kernel while the global parametric regressions are the same as shown in Figure 3.15. 95% robust confidence intervals are shown. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.17. Proportion of Patients Without Any Pharmaceutical Claims Post-ER Visit RD plots

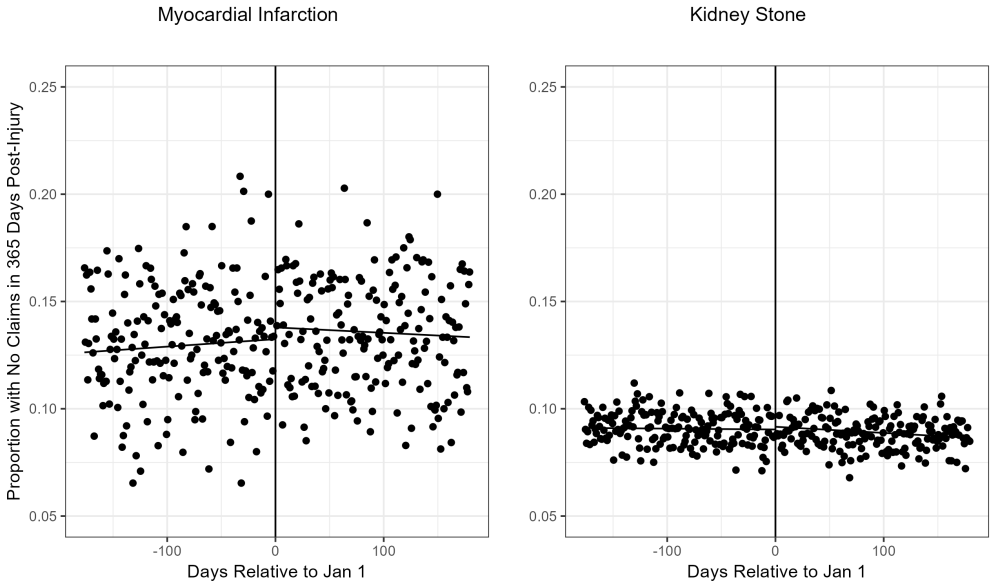

Notes: The proportion of patients who fill no claims at all in the 365 days following their myocardial infarction ER visit. There are three broad reasons why this may be: (1) they use a pharmacy not captured in the IQVIA data sample, (2) they are not adherent to their medication or are not prescribed any, (3) they may have died. Kidney Stones are shown as a comparison to evaluate the likely size of the third reason. No statistically significant discontinuities are found at the cut-point. Source: IQVIA Medical and Pharmacy Claims.

FIGURE 3.18. MI Average Coinsurance Rate Regression Discontinuity Coefficients for Those Who Have Been Sent Their Bills

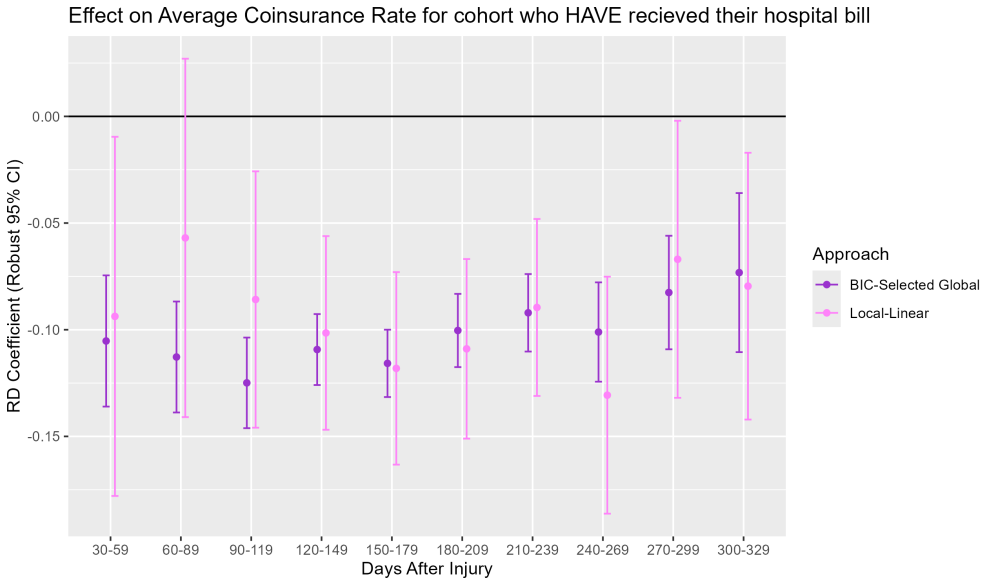

Notes: Coefficients from regression discontinuity estimates of Equations 1 and 2. The vertical axis represents the estimated difference at the cut-point in average coinsurance rate for each thirty-day period. This analysis is restricted to patients who have been sent their bills by the first day of the 30-day period following their ER visit. 95% robust confidence intervals are shown. Compare results to Figure 3.10 which includes all patients regardless of bill timing. Source: IQVIA Medical and Pharmacy Claims.

CHAPTER 4

The Preliminary Impact of the Affordable Care Act's Maximum Out-of-Pocket Limit (co-authored by Rory Martin)

4.1. Introduction

The Affordable Care Act (ACA) was a healthcare reform law enacted in March 2010 with the aim of making health insurance more affordable, expanding Medicaid, and supporting innovative models for care delivery. Although much has been written about the overall changes due to the ACA, little has been published about the impact of one of its key provisions: the requirement of compliant health plans to limit cost sharing for beneficiaries by setting an annual maximum for out-of-pocket (OOP) costs. This study analyzed the impact of ACA's maximum OOP limit on privately insured patients and found it lowered patient OOP costs and increased patient drug utilization.

Starting in 2014, the ACA established a maximum for OOP limits for private insurance plans. When insured patients reach an OOP limit their treatment costs become zero, which helps protect patients with severe medical conditions from high costs. Although the ACA's maximum OOP limit applied to all private plans, at first glance it appears many beneficiaries would have been unlikely to be affected by it. For example, in 2013 73% of covered workers were already enrolled in a plan with an OOP limit of $5,999 or less (Kaiser Family Foundation 2017b), which was a lower, more stringent limit than the ACA's $6,350 limit. Furthermore, fewer than 1% of private group insurance beneficiaries had OOP costs higher than $6,350 in 2013 (Glied and Zhu 2020). Also, policymakers had previously implemented maximum OOP limit legislation for government sponsored plans: Medicaid has had an OOP limit of

5% of family income for decades, and in 2011 Medicare Advantage plans were required to have an OOP limit for services covered under Medicare Part A and B, but not for Part D. Nonetheless, an additional component of the ACA provisions was that the OOP limit applied jointly to medical and pharmacy expenditure, and this part of the legislation was more likely to have a wider impact on patients. For example, in 2009, 85% of beneficiaries in PPOs with an OOP limit had plans that didn't count prescription drug spend towards meeting the OOP limit (Kaiser Family Foundation 2020).

4.2. How the ACA maximum OOP limit was implemented

Due to potential technical challenges insurers may have faced when combining their medical and pharmacy systems, regulators gave insurance companies a "safe-harbor" year in 2014, which delayed this part of the regulation to give insurance companies time to comply. The ACA's maximum OOP limit was implemented in two main phases for private insurance: separate maximums for pharmacy and medical benefits in 2014, and a single, combined maximum in January 2015, as illustrated in Figure 4.1. Each year, the maximum OOP limit is set by The Department of Health and Human Services to account for inflation in healthcare costs.

To determine the impact of the maximum OOP limit, three time periods were studied: 2012-13, prior to the maximum OOP limit; the "safe-harbor" phase in 2014; and the second phase from 2015 to 2018 when the combined pharmacy and medical maximum for OOP limits was implemented. For the current study, the impact of these changes on the pharmaceutical market was measured using the proportion of prescriptions paid for by private insurance that were purchased with a zero-dollar copay. This statistic is a salient outcome that affects insurers, manufacturers, and patients alike: "free to the patient" prescriptions. In 2012-13, only 15.9% of branded prescription volume corresponded to a zero-dollar copay, which rose to 19.5% in 2014 and to 25.3% in 2016, the second year of the combined maximum for OOP

limits, as shown in Figure 4.2. From 2013 to 2016, the volume of branded scripts with a zero-dollar copay saw a relative increase of 59% and an absolute increase of almost 10 percentage points. Details of the analysis are provided in the Analysis Methods and Data section.

Given that prior to 2014, fewer than 1% of covered workers had more than $6,350 in OOP costs and 73% already had an OOP limit, it is surprising to see a 4-percentage point increase in zero-dollar scripts between 2013 and 2014. However, patients who are prescribed and fill branded medications face much higher healthcare costs than the average covered worker. Because of this, the direct effect of this portion of the ACA on the branded pharmaceutical market was larger than may have been expected.

The second phase beginning in 2015 also had a substantial impact: by 2016, when the rate appears to have stabilized, the proportion of scripts with no OOP costs had increased an additional 6 percentage points. Additional longitudinal patient analysis confirmed that almost all of this increase in zero-dollar branded prescriptions was due to patients reaching their OOP limit more often and earlier in the plan year (results not shown).

The impact of the maximum OOP limit varied substantially by therapeutic area, with largest impact in therapeutic areas where patients had high cost exposure. For instance, in the multiple sclerosis market, zero-dollar copay prescriptions increased 18 percentage points with most of the impact taking place between 2013 and 2014, as illustrated in Figure 4.3. This may be due to the high cost of multiple sclerosis pharmaceutical products. On the other hand, for the diabetes market, more of the impact occurred between 2014 and 2015 potentially due to high medical costs associated with diabetes and its co-morbidities (the American Diabetes Association has estimated that almost half of the $16,752 in average medical expenditure per year for diabetes patients comes physician office visits and inpatient care (American Diabetes Association 2018)). Meanwhile some therapeutic areas, like ADHD, were not as heavily affected, possibly because ADHD patients are less likely to have high OOP costs in other disease areas.

Another way of quantifying the impact of the maximum OOP limit is to study the proportion of patients hitting an OOP limit. Each year, this proportion starts at zero in January and increases until it reaches a peak at the end of December, as shown in Figure 4.4. In 2012, two years before the ACA came into effect, only about 5% of patients reached an OOP limit (of course, some patients had no OOP limit at all). In 2018, several years after the ACA was implemented, this figure had tripled to almost 15% of patients hitting an OOP limit.

Each calendar year there is an increase in the percent of patients with $0 copays, followed by a drop in January when the vast majority of private plans reset. This became more pronounced after the ACA changes of 2014 and 2015 and is market wide. It has become an integral part of the healthcare landscape for patients, manufacturers, and insurers.

Almost three times as many patients now reach their OOP limit each plan-year as a result of the ACA's implementation, thus it is more important than ever to understand how these patients behave when they reach this limit. Although basic economic theory predicts patients will use more pharmaceutical products when prices decrease, this is a dynamic decision-making process for patients with doctors, pharmacies, and insurers acting as intermediaries, so the question is worth careful analysis. Using a fixed-effects regression, the data showed that patients increased branded drug utilization by 11.8% in the last month of the year after reaching their OOP limit, controlling for how much utilization they had in the earlier part of the year (Figure 4.5). This raises the question of whether there was a corresponding decrease in the following January when prices return to standard insurance pricing. That is, were patients simply stockpiling? The data showed a much smaller decrease in the following January, indicating that overall utilization was increasing as a result of the OOP limit rather than patients stockpiling free product in December to use in January when their plan-year reset.

The data also showed patients increased their utilization of generic medications by 7.2% in December after reaching their OOP limit (Figure 4.5), an increase that was smaller than for branded products potentially due to generic medications being cheaper than branded drugs (at full price and in most benefit designs).

One thing that health industry insiders often forget is how complicated and obfuscated the American healthcare system is for consumers, and studies have shown that consumers do not make rational choices in healthcare (Bhargava et al. 2017). An additional analysis was performed to test whether patients who reach their OOP limit in multiple years learned from the experience, but the data showed no statistically significant evidence of such behavior.

Since we have shown that the ACA increased the proportion of patients reaching their OOP limit and that once this happened patients increased utilization, we wanted to understand the impact of this provision on total pharmaceutical utilization. We estimate that the branded prescription utilization in December increased by 0.75% due to the ACA's maximum OOP limit. If we extrapolate the behavior of December patients to the entire year after accounting for the rate at which patients reach their OOP limit throughout the year, we calculate a 0.29% increase in total branded utilization as a result of this single provision in the ACA. In 2018, this would have represented an increase in pharmaceutical gross revenue of about $1.4 billion based on our analysis of IQVIA's DDD Subnational Sales database.

The finding that patients increase utilization of both branded and generic products upon reaching their OOP limit raises the question of whether patient socioeconomic status, which may be associated with being better informed, plays a role. To explore this further, the impact of patient income on drug utilization was tested. No direct measure of patient income was available, so patients were segmented based on the income quartile of the Census Public Use of Microdata Area (PUMA) and ZIP code geographical areas in which they lived.[1]

[1]The PUMA geographical area provided the advantage of being able to observe the income for privately-insured individuals separately from uninsured, Medicaid, and Medicare populations. However, PUMAs are a coarse geographic unit and one PUMA may cover many neighborhoods of differing income levels. Meanwhile, ZIP codes are much more granular geographic areas, but income data for privately-insured individuals was

Patients in the lowest income quartile geography increased their utilization four times as much as patients in the highest quartile geography once they reached their OOP limit, 20% versus 5%, respectively (see Figures 4.6 and 4.7). This is evidence that lower income patients benefit the most from a maximum OOP limit and that increased price sensitivity among lower income patients outweighs any information advantage that higher income patients may have. The analysis also found evidence that wealthier patients engage in product warehousing, advancing purchases from January into December. This is consistent with the notion that low-income patients struggle to pay for drugs and increase overall utilization when it is cheaper to do so, while wealthier patients are simply reducing OOP costs while maintaining their existing medication schedule. Encouragingly, this analysis is evidence that the OOP limit has substantially helped low income patients who have experienced very high OOP costs in the calendar year to fill their prescriptions at the end of the year.

4.3. Conclusion

Since the ACA was enacted in 2010 it has faced a number of legal challenges. Additionally, in 2019 the Tax Cuts and Jobs Act of 2017 rescinded the federal tax penalty for violating the individual mandate of the ACA. The current paper has for the first time estimated the impact of the ACA's maximum OOP limit, quantifying what would be at risk if this part of the ACA were to be repealed or otherwise compromised.

The current study found that branded zero-dollar copay commercial prescriptions increased by 4 percentage points between 2013 and 2014 upon the initial implementation of the law, and by an additional 6 percentage points between 2014 and 2016 upon the combining of medical and pharmacy expenses. However, the impact varied substantially by therapeutic area. Disease areas with expensive drug products and/or high-cost patient populations saw larger effects than those with less severe co-morbidities and lower costs.

---

not available at that level. Median household income for the whole ZIP code was used instead to construct the income quartiles.

The introduction of the ACA's maximum OOP limit likely accelerated the use of deductibles for pharmacy benefits. For example, the Kaiser Family Foundation's annual survey of employer health benefits found deductibles rose four times faster than premiums in 2016 (Kaiser Family Foundation 2016) and have continued to rise ever since. This has had multiple effects. First, by shifting costs from beneficiaries with the highest treatment expenses to beneficiaries in general, it may have contributed to the observed increase in the percent of beneficiaries who reached their OOP limit. Second, it increased patient awareness of the true cost of pharmaceutical products because during their deductible period, beneficiaries are exposed to the full cost of those products.

The OOP limit, along with the high prevalence of large and increasing deductibles, are shaping the commercial pharmaceutical landscape and lead to an annual cycle of high patient OOP costs in the beginning of the calendar year and lower OOP costs as the end of the year approaches. This annual cycle underpins the strategy of many market participants. For example, it incentivizes pharmaceutical manufacturers to use coupons to help patients stay on therapy through the expensive, early part of the year in the hope that they will still be adherent to therapy once deductibles have faded and OOP limits have been reached. In response, pharmacy benefit managers (PBMs) and insurers are incentivized to ensure coupon payments do not count toward deductibles or OOP limits through accumulator programs to maintain utilization control of branded pharmaceuticals. Finally, patients increase their utilization after reaching their OOP limit.

The current study looked at the impact of the ACA's maximum OOP limit on patient OOP costs and drug utilization. It's also of interest to know how manufacturers and insurers responded to this legislation. For example, did manufacturers change the timing or magnitude of price increases, and if so, how did insurers respond? Such questions may be the basis of follow-up studies from the current authors.

In the current study, drug utilization of patients increased 7-12% once they reached their OOP limit for the year. The effect was higher for branded products than generics and was four times larger in low income geographies than high income ones. On the question of whether patients are likely to stockpile product at the end of the year and then reduce purchases in the beginning of the year once prices reset, the data found suggestions of that behavior among the highest income geographies, but the predominant observed behavior was an increase in pharmaceutical utilization in December without a commensurate decrease in January. These findings support previous studies that indicate that high out-of-pocket costs decrease adherence to medication (Doshi et al. 2009).

As the market evolves and policy changes, many factors will help shape the impact of the ACA's maximum OOP limit in the future. If cost-sharing and prices increase faster than statutory OOP limits, the annual cycle of high prices in the beginning of the year and low prices at year's end will grow stronger. However, if patients shift out of ACA-compliant healthcare plans in large numbers, as might happen after the repeal of the insurance mandate, the maximum OOP limit will be less impactful for market participants. In the meanwhile, policymakers and market participants should be aware of the many effects that the ACA's maximum OOP limit has had on patient cost-sharing in the U.S. healthcare system.

4.4. Analysis Methods and Data

To measure the impact of the maximum OOP limit, prescription volume was estimated across all pharmaceutical products found in IQVIA's Longitudinal Access and Adjudication Dataset (LAAD) reference data. This spans all U.S. pharmaceutical products, including branded and generic products, patient- and physician-administered products, and all disease areas. Because scripts can be written for different quantities of medication, script volume was weighted by days of supply. Thus, a script for a 90-day supply of a drug would have three times the weight of a 30-day supply script.

Regarding the proportion of patients reaching their OOP limit (Section 1), most but not all patients have plan years that align with the calendar year, meaning their deductibles and OOP limits reset in January.1 Data indicating which patients were exceptions to this was not available, thus this was not taken into account in the analysis in Figure 4.4.

The impact on patient behavior due to a patient reaching his or her OOP limit was estimated as follows. In a year in which their OOP limit was reached, the change in drug utilization was measured: December utilization versus pre-December utilization. The same quantity was estimated in a year in which the patient did not reach his or her OOP limit, to establish a baseline, and the difference was calculated. The process was repeated for drug utilization in January to test for the presence of product warehousing, i.e., flat total utilization where January purchases are advanced into December. Except where indicated, the current study looked only at branded prescriptions since generic prescriptions often have a zero OOP cost due to benefit design.

The following regression was estimated separately for the last four weeks of the calendar year (December) and the first four weeks of the year (January):

$$y_{it} = \alpha_t + \delta_i + \beta_1 x_{it} + \gamma_1 z_{it}^B + \gamma_2 z_{it}^G$$

Where y_{it} is the number of days of therapy in the last (first) four weeks of the calendar year and x_{it} is an indicator whether or not the patient had reached the MOOP limit by the last four weeks of the calendar year. z_{it}^B and z_{it}^G are the number of days of therapy in the surrounding 44 weeks (July-December and late January to June) for branded and generic medications respectively. Lastly, α_t and δ_i are time and individual fixed effects. β_1 is the coefficient of interest which measures the increase in days of therapy associated with reaching an OOP limit controlling for variables related to the patient, year, and pharmaceutical utilization outside of the critical periods near the end and beginning of the calendar year.

A handful of confounding effects existed in the pharmaceutical market around the time the ACA was enacted, including coupon usage, accumulator programs[2], an increased presence of generics, and price increases for branded products. The current study measured patient OOP costs before the use of coupons took effect. Unless an accumulator program was being used, a coupon would reduce OOP costs for the patient but wouldn't impact the patient reaching his or her OOP limit. An increased generic presence may have caused fewer patients to hit their OOP limit than otherwise would have done, while price increases would have had the opposite effect. The current study did not control for generics and price increases because its aim was to measure the impact of the ACA's maximum OOP limit on market conditions as they existed at the time the ACA was enacted.

[2]Accumulator programs are a strategy used by insurers PBMs to prevent manufacturer-provided coupons from accumulating to the deductible threshold thereby preventing manufacturers from "buying patients through their deductible phase". They were primarily introduced in 2018, after the data used in this study.

4.5. Figures

FIGURE 4.1. How the ACA's maximum OOP limit was implemented

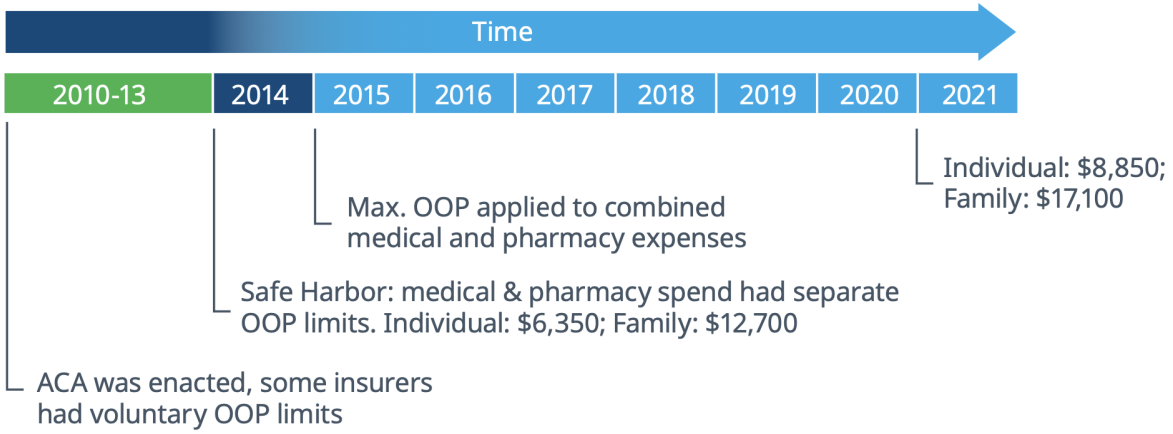

FIGURE 4.2. Percent of branded prescriptions with a zero-dollar copay dispensed to privately insured patients

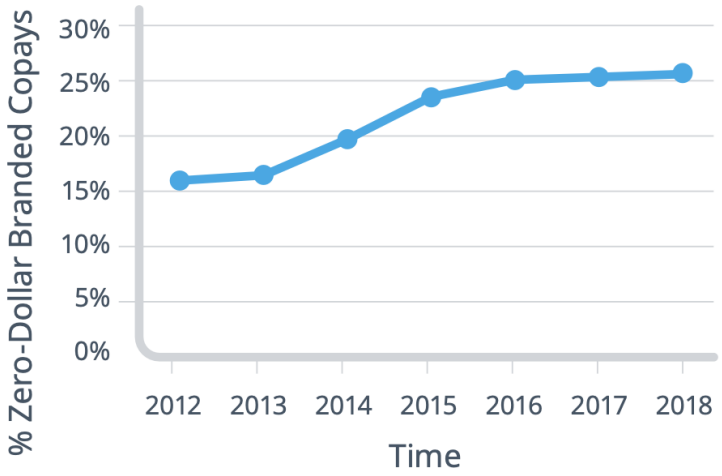

FIGURE 4.3. Proportion of zero-dollar copay branded prescriptions by therapeutic area and year

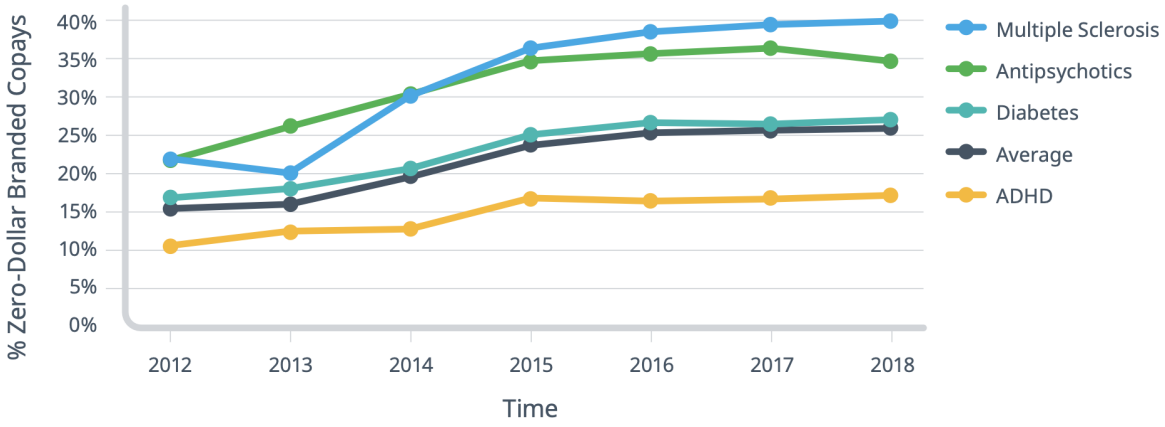

FIGURE 4.4. Percent of patients reaching their OOP limit

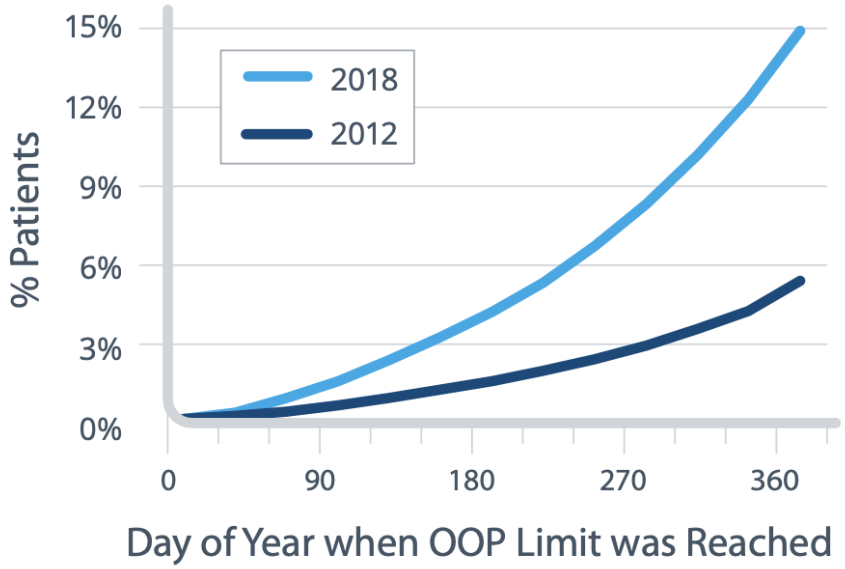

FIGURE 4.5. Impact of reaching an OOP limit on therapy purchased in December and January

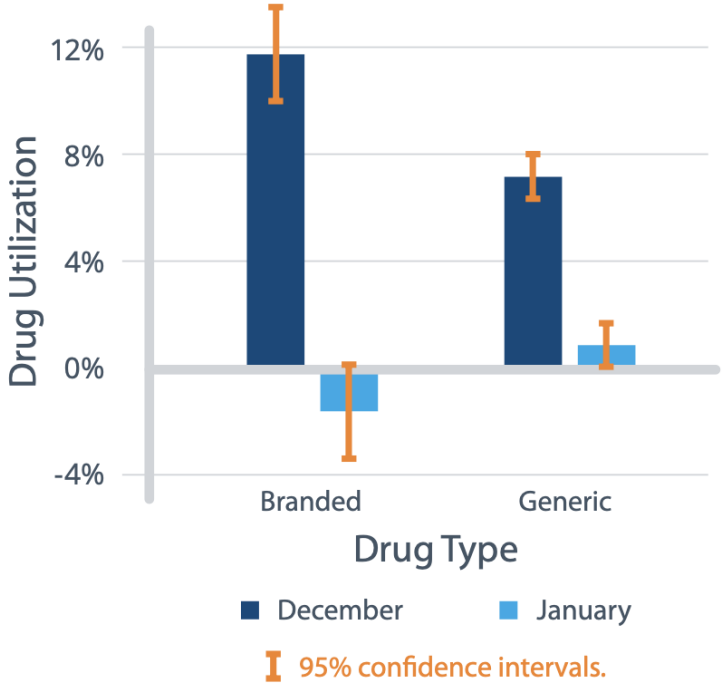

FIGURE 4.6. Impact of reaching an OOP limit on branded products purchased in December and January by census PUMA income quartile geography

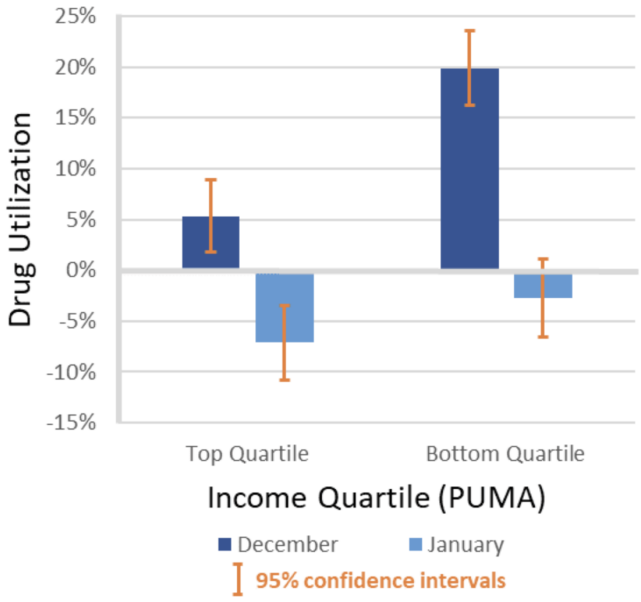

FIGURE 4.7. Impact of reaching an OOP limit on branded products purchased in December and January by ZIP code income quartile geography

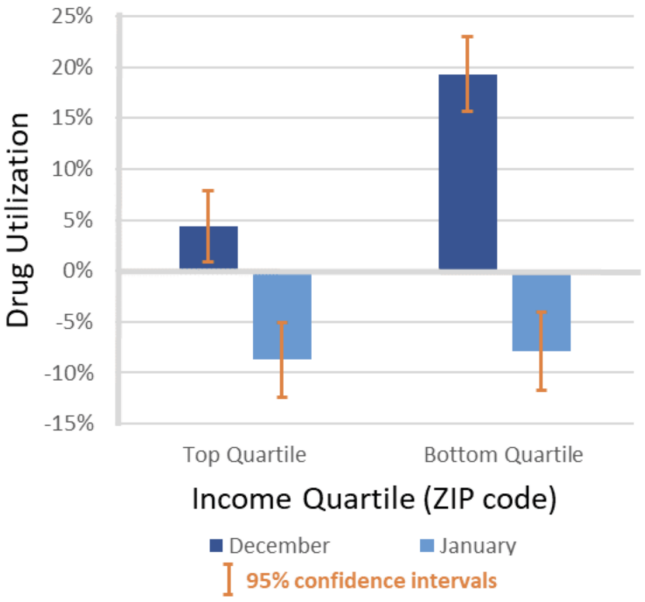

Bibliography

Abaluck, Jason and Jonathan Gruber, "Choice inconsistencies among the elderly: evidence from plan choice in the Medicare Part D program," *American Economic Review*, 2011, *101* (4), 1180–1210.

American Diabetes Association, "Economic costs of diabetes in the US in 2017," *Diabetes care*, 2018, *41* (5), 917–928.

Angrist, Joshua D and Guido W Imbens, "Two-stage least squares estimation of average causal effects in models with variable treatment intensity," *Journal of the American statistical Association*, 1995, *90* (430), 431–442.

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen, "Moral hazard in health insurance: Do dynamic incentives matter?," *Review of Economics and Statistics*, 2015, *97* (4), 725–741.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, "Choose to lose: Health plan choices from a menu with dominated option," *The Quarterly Journal of Economics*, 2017, *132* (3), 1319–1372.

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," *The Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna, "Difference-in-differences with a continuous treatment," *arXiv preprint arXiv:2107.02637*, 2021.

Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell, "Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs," *The Econometrics Journal*, 2020, *23* (2), 192–210.

Card, David, Carlos Dobkin, and Nicole Maestas, "Does Medicare save lives?," *The quarterly journal of economics*, 2009, *124* (2), 597–636.

Chandra, Amitabh, Evan Flack, and Ziad Obermeyer, "The health costs of cost-sharing," *The Quarterly Journal of Economics*, 2024, p. qjae015.

Curkendall, S, V Patel, M Gleeson, RS Campbell, M Zagari, and R Dubois, "Compliance with biologic therapies for rheumatoid arthritis: do patient out-of-pocket payments matter?," *Arthritis Care & Research: Official Journal of the American College of Rheumatology*, 2008, *59* (10), 1519–1526.

Cutler, David M and Richard J Zeckhauser, "The anatomy of health insurance," in "Handbook of health economics," Vol. 1, Elsevier, 2000, pp. 563–643.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo, "The economic consequences of hospital admissions," *American Economic Review*, 2018, *108* (2), 308–52.

Doshi, Jalpa A, Jingsan Zhu, Bruce Y Lee, Stephen E Kimmel, and Kevin G Volpp, "Impact of a prescription copayment increase on lipid-lowering medication adherence in veterans," *Circulation*, 2009, *119* (3), 390–397.

Einav, Liran and Amy Finkelstein, "Moral hazard in health insurance: what we know and how we know it," *Journal of the European Economic Association*, 2018, *16* (4), 957–982.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P Newhouse, Heidi Allen, Katherine Baicker, and the Oregon Health Study Group, "The Oregon health insurance experiment: evidence from the first year," *The Quarterly journal of economics*, 2012, *127* (3), 1057–1106.

Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo, "Hospitals as insurers of last resort," *American Economic Journal: Applied Economics*, 2018, *10* (1), 1–39.

Glied, Sherry A and Benjamin Zhu, "Catastrophic out-of-pocket health care costs: a problem mainly for middle-income Americans with employer coverage," *Commonwealth Fund*, 2020.

Goldman, Dana P, Geoffrey F Joyce, Jose J Escarce, Jennifer E Pace, Matthew D Solomon, Marianne Laouri, Pamela B Landsman, and Steven M Teutsch, "Pharmacy benefits and the use of drugs by the chronically ill," *Jama*, 2004, *291* (19), 2344–2350.

Gruber, Jonathan and Benjamin Sommers, "The Affordable Care Act?s Effects on Patients, Providers and the Economy: What We?ve Learned So Far," Technical Report w25932, National Bureau of Economic Research, Cambridge, MA June 2019.

Hackney, Donald D, Daniel Friesner, and Erica H Johnson, "What is the actual prevalence of medical bankruptcies?," *International Journal of Social Economics*, 2016.

Hausman, Catherine and David S Rapson, "Regression discontinuity in time: Considerations for empirical applications," *Annual Review of Resource Economics*, 2018, *10* (1), 533–552.

Hettinger, Katlyn, "Intertemporal Substitution in Response to Non-Linear Health Insurance Contracts." PhD dissertation, Michigan State University 2023.

HHS, "Affordable Care Act Implementation FAQs - Set 12 Q2," 2013.

_____, "Patient Protection and Affordable Care Act; HHS Notice of Benefit and Payment Parameters for 2016," Technical Report 2015.

Himmelstein, David U, Deborah Thorne, Elizabeth Warren, and Steffie Woolhandler, "Medical bankruptcy in the United States, 2007: results of a national study," *The American journal of medicine*, 2009, *122* (8), 741–746.

Hung, Anna, Dan V Blalock, Julie Miller, Jaime McDermott, Hannah Wessler, Megan M Oakes, Shelby D Reed, Hayden B Bosworth, and Leah L Zullig, "Impact of financial medication assistance on medication adherence: a systematic review," *Journal of managed care & specialty pharmacy*, 2021, *27* (7), 924–935.

Huskamp, Haiden A, Patricia A Deverka, Arnold M Epstein, Robert S Epstein, Kimberly A McGuigan, and Richard G Frank, "The effect of incentive-based formularies on prescription-drug utilization and spending," *New England Journal of Medicine*, 2003, *349* (23), 2224–2232.

Illich, Kepler, "Free Drugs in December: The Effects of the Affordable Care Act's Maximum Out-of-Pocket Limit Regulation." PhD dissertation, University of California, Davis November 2024.

IQVIA Institute, "Available IQVIA Data," July 2022.

Kaiser Family Foundation, "2013 Employer Health Benefits Survey," 2013.

_____, "2015 Employer Health Benefits Survey," 2015.

_____, "The Missing Debate Over Rising Health-Care Deductibles," Technical Report, Kaiser Family Foundation September 2016.

_____, "2017 Employer Health Benefits Survey," Technical Report 2017.

_____, "Employer Health Benefits Survey," Technical Report, Kaiser Family Foundation 2017.

_____, "Increases in Cost-Sharing Payments Continue to Outpace Wage Growth," 2018.

_____, "2019 Employer Health Benefits Survey," 2019.

_____, "Potential Impact of California v. Texas Decision on Key Provisions of the Affordable Care Act," Technical Report, Kaiser Family Foundation September 2020.

_____, "Deductibles in ACA Marketplace Plans, 2014-2024," Technical Report, Kaiser Family Foundation 12 2023.

_____, "Employer Health Benefits Survey," Technical Report, Kaiser Family Foundation 11 2023.

Karter, Andrew J, Melissa M Parker, Matthew D Solomon, Courtney R Lyles, Alyce S Adams, Howard H Moffet, and Mary E Reed, "Effect of out-of-pocket cost on medication initiation, adherence, and persistence among patients with type 2 diabetes: The Diabetes Study of Northern California (DISTANCE)," *Health services research*, 2018, *53* (2), 1227–1247.

Keisler-Starkey, Katherine, Lisa N Bunch et al., "Health insurance coverage in the United States: 2020," *Washington, DC: US Census Bureau*, 2021.

Kirkman, M Sue, Megan T Rowan-Martin, Rebecca Levin, Vivian A Fonseca, Julie A Schmittdiel, William H Herman, and Ronald E Aubert, "Determinants of adherence to diabetes medications: findings from a large pharmacy claims database," *Diabetes care*, 2015, *38* (4), 604–609.

Kowalski, Amanda, "Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care," *Journal of Business & Economic Statistics*, 2016, *34* (1), 107–117.

Lee, David S and David Card, "Regression discontinuity inference with specification error," *Journal of Econometrics*, 2008, *142* (2), 655–674.

_____ **and Thomas Lemieux**, "Regression discontinuity designs in economics," *Journal of economic literature*, 2010, *48* (2), 281–355.

Manning, Willard G, Joseph P Newhouse, Naihua Duan, Emmett B Keeler, and Arleen Leibowitz, "Health insurance and the demand for medical care: evidence from a randomized experiment," *The American economic review*, 1987, pp. 251–277.

McCrary, Justin, "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of econometrics*, 2008, *142* (2), 698–714.

McNamara, Robert L, Kevin F Kennedy, David J Cohen, Deborah B Diercks, Mauro Moscucci, Stephen Ramee, Tracy Y Wang, Traci Connolly, and John A Spertus, "Predicting in-hospital mortality in patients with acute myocardial infarction," *Journal of the American College of Cardiology*, 2016, *68* (6), 626–635.

Neumann, Peter J, Eileen A Sandberg, Chaim M Bell, Patricia W Stone, and Richard H Chapman, "Are Pharmaceuticals Cost-Effective? A Review Of The Evidence: Do drug treatments give value for the money? Careful analysis can yield useful information, this study finds.," *Health affairs*, 2000, *19* (2), 92–109.

Newhouse, Joseph P, *Free for all?: lessons from the RAND health insurance experiment*, Harvard University Press, 1993.

Shrank, William H, Niteesh K Choudhry, Joshua N Liberman, and Troyen A Brennan, "The use of generic drugs in prevention of chronic disease is far more cost-effective than thought, and may save money," *Health affairs*, 2011, *30* (7), 1351–1357.