

Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets*

Jason Abaluck, Yale and NBER
Jonathan Gruber, MIT and NBER
Ashley Swanson, University of Pennsylvania and NBER

April 25, 2016

Abstract

Medicare Part D enrollees face a complicated decision: they dynamically choose prescription drug consumption in each period given difficult-to-find prices and a non-linear budget set. We use Part D claims data to estimate a flexible model of consumption that accounts for non-linear prices, dynamic responses, and salience. We use reduced form price responses from a linear regression of consumption on coverage range prices to compare performance under several models of behavior. We find small price elasticities, substantial myopia, and that salient characteristics impact consumption beyond their effect on prices. A hyperbolic discounting model with salience fits the data best.

Keywords: Moral hazard, nonlinear budget sets, salience, Medicare, prescription drugs

1 Introduction

Under the Medicare Part D prescription drug benefit, private insurers offer a wide range of products with varying prices and features. While the government sets a standard insurance design, over 90% of enrollees are enrolled in non-standard plans, subject to the constraint that these alternatives be at least as generous as the standard plan. In particular, insurance plans offered through Medicare Part D vary widely in their deductibles, in the copayments

*We are grateful to Zack Cooper, Christina Dalton, Liran Einav, Amy Finkelstein, Gautam Gowrisankaran, Jerry Hausman, Kyoungae Jung, Amanda Kowalski, Fiona Scott Morton, Robert Town, and numerous seminar participants for helpful comments. Kathleen Easterbrook, Ayesha Mahmud, and Adrienne Sabety provided outstanding research assistance. Financial support was provided by the National Institute of Aging.

and coinsurance for prescriptions above the deductible, and in their coverage of drugs in the infamous Part D “donut hole” where the standard plan offers no coverage. A number of analysts have expressed concern both about the overall generosity of the standard package, and about consumer confusion in choosing across this wide variety of alternatives. Responding to the former concern, the Patient Protection and Affordable Care Act (ACA) of 2010 “fills in” the donut hole for the standard plan design.

The welfare impacts of these variations in plan design, as well as of filling in the donut hole in the future, depend critically on how prescription drug spending responds to Part D coverage. But the complex nature of Part D contracts makes it difficult to correctly model the effects of Part D coverage on drug spending. Enrollees face a complicated non-linear budget constraint for their drug consumption decisions, whereby both current and future prices for drugs are a function of consumption to date, as well as a dynamic environment in which uncertainty is realized gradually over time. The complexity of this optimization problem may be particularly onerous for the elderly Part D population.

In this paper, we present a model of prescription drug coverage under Part D which accounts for optimization with non-linear budget sets, dynamic incentives due to uncertainty and myopia, and variation in salience of different aspects of the insurance contract. We use flexible reduced form regressions to determine how enrollees respond to price changes throughout the budget set and relate these to several structural models of consumption behavior. Notably, our identification of responses to price *changes* will be appropriate for modeling counterfactuals such as filling in the donut hole, in that the vast majority of enrollees at any given time were enrolled in the same plan in the previous year; enrollees may be less attentive to changes in the coverage characteristics of their existing plans than new enrollees would be to differences in the coverage characteristics of alternative plans.

The standard approach to estimating demand in the presence of non-linear budget sets is to estimate a nonlinear structural model assuming a particular model of optimization behavior as in Hausman (1985) and Kowalski (2015). However, traditional models of a

response to nonlinear budget sets assume that all price responses reflect a single underlying parameter that determines price elasticities. More recent examples relax this assumption to incorporate behavioral responses such as myopia, typically modeled using a discount factor; see Einav, Finkelstein, and Schrimpf (2015) and Dalton, Gowrisankaran, and Town (2015) for two excellent recent examples in the Part D context. A central element of these newer approaches is to achieve identification from the response of individuals near “kinks” in the budget constraint, such as the one created by the donut hole in Part D plans.

These approaches face three key limitations. First, as has been discussed in detail in the energy economics literature, lack of information or understanding about prices may lead consumers in nonlinear contracts (such as insurance contracts and electricity contracts) to use rules of thumb beyond discounting in determining their consumption. For example, Liebman and Zeckhauser (2004) note that individuals may respond to nonlinear price schedules by “ironing” (responding to average or local average prices), or by “spotlighting” (responding to immediate “spot” prices rather than to future marginal prices). E.g., Brot-Goldberg et al. (2015) find that individuals in nonlinear health insurance contracts respond only to spot prices, even when there is very little uncertainty regarding end of year price. “Spotlighting” behavior can be captured in the more recent papers cited above; “ironing” can not.

Second, just as consumers may not be perfectly forward-looking in their consumption behavior, they may also be confused about how visible changes in benefit coverage impact the prices they face. A large recent literature in economics highlights the role of price “salience,” which leads consumers in complex or nontransparent environments to be inattentive to some prices (see Chetty, Looney, and Kroft (2009) for a review). Behavioral biases in consumption may impact consumer weighting of particular coverage characteristics based on visibility. Empirical studies of taxation have found evidence of such biases; for example, Feldman, Katuscak, and Kawano (2015) find that consumers responded to the removal of the lump sum Child Tax Credit by reducing their labor income, even though the removal did not impact their net wages. The structural empirical literature in health care has not typically

allowed consumer responses to vary with price salience.

Third, most individuals who consume prescription drugs are not near the budget constraint kink created by the donut hole. Indeed, only 10% of individuals in a given year end the year within \$200 of the donut hole, and that 10% is observably sicker than the average enrollee, with higher medical spending and more chronic illnesses. Moreover, many individuals who are impacted by policy counterfactuals such as filling in the donut hole are those who end the year well past the donut threshold – according to our estimates, only 14% of the effect of filling in the donut hole is accounted for by enrollees ending the year within \$200 of the donut hole kink. It is therefore unclear whether the results identified by this particular variation extend more broadly to the population of Part D enrollees.

We present a new model of prescription drug demand that addresses all of these concerns. To do so we draw on the spirit of the nonparametric estimation approach in Blomquist and Newey (2002), which instead of imposing one structure that incorporates all budget segments (e.g., that a particular level of myopia completely explains the relative responses to current and future prices) allows the data to tell us how enrollees’ consumption throughout the year responds to different budget set segment prices. We also allow enrollees to respond to variables capturing nominally large changes in benefit coverage that may be more salient. This reduced form model is then embedded in a structural model of drug consumption which allows us to better assess which beneficiaries are impacted by changing prices and how these prices impact the time path of consumption.

We demonstrate that linear regression methods can be used to recover parameters from structural models of consumption for individuals whose marginal prices are in the interiors of budget set segments. This approach allows us to estimate price sensitivity for individuals throughout the Part D spending distribution and not just within a relatively narrow spending range near the budget set kink.¹ One advantage of our approach relative to existing

¹Our approach involves explicitly dropping individuals from our estimation sample who are “close” to the kink. It is important to note that, for identification purposes, we are still using a greater range of identifying variation than other studies focusing on behavior near the kink for identification. Even though those studies may not drop observations well away from the kink from their estimation samples, identification in those

approaches is that we can estimate structural parameters without assuming knowledge of the full distribution of consumer uncertainty. Instead, we assume only that we can forecast with very high probability the likelihood that consumers will end the year in a particular portion of the budget set. A concern with this approach is that consumers may still be responding to their low probability of bunching or switching. To probe robustness, we estimate a version of our model which explicitly allows for uncertainty and endogenous coverage phase transitions using conventional methods and find that our results are largely unchanged. We also use our estimates to simulate nonlinear consumption behavior: we can test whether price responses for individuals in the estimation sample can be generalized to the (much smaller set of) individuals near the budget set kink.

The main data source is a 20% sample of Medicare Part D claims provided by the Centers for Medicare and Medicaid Services (CMS). The claims data include information on drugs consumed, as well as the date, quantity, wholesale price and amount paid by insurer and beneficiary for each claim. Our identification strategy utilizes the significant year-to-year variation in the cost-sharing features of Part D plans for existing Part D enrollees. While new enrollees might be expected to pay more attention to prices, 90% of Part D enrollees remain in a given plan across years, so that the response of these consumers is most predictive for changes in the structure of Part D, such as filling in the donut hole. For this counterfactual and others impacting existing (inertial) enrollees in health insurance contracts, this strategy is more representative than others in the literature that consider price responses to the introduction of insurance.²

To illustrate the thought experiment behind our identification, we begin our analysis with a differences-in-differences analysis of plans filling in the donut hole. The results show that enrollees' consumption is sensitive to the presence of donut hole coverage, and that enrollees' responses increase in magnitude as more individuals enter the donut hole.

cases comes from the small fraction of beneficiaries who bunch.

²See, e.g., Duggan, Healy, and Scott Morton (2008) regarding the response of new enrollees when Part D was introduced.

We extend this analysis by estimating reduced form regressions of year-to-year consumption changes on changes in the price along each budget segment and changes in salient coverage characteristics. Our results offer mixed support for previous approaches. We find that consumers’ responses to different coverage phase prices vary steeply with the proportion of enrollees currently in those coverage phases, even holding marginal coverage phase fixed. We show that a model of myopia similar to Aron-Dine, et al. (2015) fits this pattern well. Notably, the preferred structural model performs quite well in capturing consumption patterns not only in the estimation sample but also in predicting nonlinear consumption behavior for individuals near the kink (i.e., “bunching”). On the other hand, we find evidence of substantial price “salience.” In particular, we find that plan donut coverage impacts consumption more than would be expected given its impact on either current or expected marginal price. Our most striking evidence of salience is that even low-spending individuals who are highly unlikely to enter the donut hole are nonetheless responsive in their consumption to the presence or absence of donut hole coverage.

Our reduced form results suggest price elasticities of around -0.13 on average, which is of a similar magnitude to the previous literature on prescription drug and health care services demand. The dynamics in the observed marginal price responses imply an estimated (quarterly) β (in a $\beta - \delta$ discounting model) of 0.31, suggesting a very high degree of myopia. But we also estimate sizable salience effects – we find for example that low spending enrollees who are unlikely to hit the donut hole reduce consumption by about 6% when they lose gap coverage. Given our elasticities, this implies that they respond to gap coverage (which they are unlikely to ever use) the same way they would respond to an 81% increase in prices.³ This very large effect is assumed to be zero in the current literature on utilization effects.

Using the structural parameters implied by our linear estimates, we simulate consumption responses over the entire nonlinear budget set and estimate the consumption response to the counterfactual of filling in the Part D donut hole. We demonstrate that, given our estimates,

³This calculation is based on a comparison of the ICR price and “Stark” coefficients in Table 6: $-0.099 * (\ln(p_2) - \ln(p_1)) = -0.059$ implies p_2 is 81% higher than p_1 .

it matters not just what prices are changed, but also when they change in the year and how the price changes are presented. Filling in the donut hole will lead to substantial increases in consumption, but such increases will be realized unevenly over the year and will affect even low-spending parts of the Part D population due to price salience. Salience effects account for over 30% of the total consumption response to donut hole coverage.

The rest of the paper proceeds as follows. In Section 2, we describe the background of the Part D program and the related literature on elderly decision-making, moral hazard in health care, and nonlinear prices. Section 3 describes our identification strategy and data, provides details on price variable construction, and presents a simple reduced form analysis illustrating the identification approach (capturing the impact of donut hole coverage on consumption). In Section 4, we outline a model of how consumption responds to prices in the presence of both myopia and salience in a dynamic setting. In Section 5, we present the reduced form empirical estimates. Section 6 translates our price coefficients into structural parameters and shows the results of our counterfactual simulations. Section 7 concludes.

2 Background and Related Literature

The Part D program passed in 2003, and was implemented in 2006 to provide, for the first time, subsidized prescription drug insurance for the elderly.⁴ The most noticeable innovation of Part D is that this new Medicare benefit is not delivered by the government, but rather by private insurers under contract with the government. Beneficiaries can choose from three types of private insurance plans for coverage of their drug expenditures. The first type are stand-alone plans called Medicare Prescription Drug Plans (PDPs) (plans that just offer prescription drug benefits). For example, in 2006, there were 1,429 total PDPs offered nationally, with most states offering about forty PDPs. The second are Medicare Advantage (MA) plans, plans that provide all Medicare benefits, including prescription drugs, such

⁴Duggan et al. (2008) provide a detailed overview of the Part D program and many of the economic issues it raises, so we just provide a brief overview here.

as HMO, PPO, or private FFS plans. There were 1,314 MA plans nationally in 2006. Finally, beneficiaries could retain their current employer/union plans, as long as coverage is “creditable” or at least as generous (i.e. actuarially equivalent) as the standard Part D plan, for which they would receive a subsidy from the government. We focus on PDP plans so that other health benefits are held constant.

Under Part D, recipients are entitled to basic coverage of prescription drugs by a plan with equal or greater actuarial value to the standard Part D plan. The standard plan for the year 2006 covers: none of the first \$250 in drug costs each year; 75% of costs for the next \$2,000 of drug spending (up to \$2,250 total); 0% of costs for the next \$3,600 of drug spending (up to \$5,100 total – the infamous “donut hole”); and 95% of costs above \$5,100 of drug spending. Coverage thresholds for the standard plan have increased in each year since first implementation of the program; the standard plan deductible and donut threshold in 2009, the last year of our sample, were \$295 and \$2,700, respectively. The government also placed restrictions on the structure of the formularies that plans could use to determine which prescription medications they would insure. In practice, the vast majority of enrollees have chosen plans with non-standard cost-sharing; over 90% of beneficiaries in 2006 were not enrolled in the standard benefit design, but rather were in plans with low or no deductibles, flat payments for covered drugs following a tiered system, or some form of coverage in the donut hole. The ACA mandates that the donut hole be “filled in” gradually by 2020. For the 2014 benefit year, enrollees in plans without coverage in the donut hole are entitled to a 52.5% discount on branded drugs and a 21% discount on generics while in the donut hole.⁵

Enrollment in Part D plans was voluntary for Medicare eligible citizens. In order to mitigate adverse selection, Medicare recipients not joining during their first eligibility period (and who did not have other creditable prescription drug coverage) were subject to a financial penalty if they eventually joined the program.⁶

⁵Drug manufacturers offer the branded discount under the Medicare Coverage Gap Discount Program; Medicare covers the 21% generic discount in the donut hole (CMS, 2015).

⁶One group was automatically enrolled: low income elders who had been receiving their prescription drug coverage through state Medicaid programs (the “dual eligibles”). These dual eligibles were enrolled in

Our project builds on several literatures. First, there is a rich literature on the impact of cost-sharing on health care utilization and this literature is reviewed in great detail in Chandra, Gruber, and McKnight (2010). Of particular note is the RAND Health Insurance Experiment (HIE), which is summarized in Manning et al. (1987). The HIE showed that consumption of medical services was modestly price responsive, with an overall estimated arc-elasticity of medical spending in the range of -0.2. A large subsequent literature has investigated utilization effects specifically in the context of prescription drugs. This literature is reviewed in Goldman, Joyce, and Zheng (2007), which finds elasticities ranging from -0.2 to -0.6. Several studies examine utilization effects specifically in the context of Medicare Part D. Lichtenberg and Sun (2007) examine the change in drug expenditures for elderly and non-elderly consumers following the introduction of Part D and estimate that Part D led to a 12.8% increase in prescription drug utilization (from an 18.4% reduction in patient cost sharing, an arc-elasticity of -0.70); Yin, et al. (2008) report a 5.9% increase in utilization in data from a large pharmacy chain. Using different data sources but a similar methodology, Ketcham and Simon (2008) estimate an arc-elasticity of -0.47. Chandra, Gruber, and McKnight (2010) analyze another group of retirees, from the California Public Employees Retirement System (CalPERS) and find an arc-elasticity of prescription drug consumption of -0.08 to -0.15. Thus, previous studies have consistently found evidence that drug utilization responds to out-of-pocket prices, but the magnitude of the estimates varies dramatically across studies. Our data include a representative sample of the entire universe of Medicare Part D claims and will thus shed light on the elasticity of demand for the full sample of unsubsidized enrollees.

Another literature on healthcare utilization models health care consumption elasticities in the presence of non-standard pricing. Kowalski (2015) studies the aggregate utilization of

Part D plans by default if they did not choose one on their own. The Part D plans for dual eligibles could charge copayments of only \$1 for generics/\$3 for name brand drugs for those below the poverty line, and only \$2 for generics/\$5 for name brand drugs for those above the poverty line, with free coverage above the out-of-pocket threshold of \$3,600. In addition, other low income groups were eligible for the Low Income Subsidy (LIS) or for other subsidy programs that lowered their premiums and cost sharing.

medical care in a non-linear budget set environment with a static consumption decision and finds consumers to have quite low price elasticities, thus concluding that generous coverage leads to modest deadweight losses from moral hazard. Aron-Dine, Einav, Finkelstein, and Cullen (2015) model dynamic consumption of medical services in the presence of a varying effective deductible and show that individuals respond not only to their expected marginal price but also to the spot price they face before reaching coverage thresholds. Einav, Finkelstein, and Schrimpf (2015) consider Part D enrollees specifically by focusing on dynamic incentives due to enrollees entering into Part D contracts at different points in the year (as they age into Medicare) and estimate an overall price elasticity from the degree of bunching observed among individuals whose total drug expenditures place them near the donut hole threshold at the end of the year. Using variation in dynamic incentives for enrollees aging into the Medicare benefit, they estimate a weekly δ of 0.96, which translates roughly to a quarterly β of 0.5; they find static price elasticities ranging from -0.3 to -0.5. In contrast, Dalton, Gowrisankaran, and Town (2015) estimate dynamic preferences based on consumption changes as individuals predictably cross the donut threshold and price elasticities based on cross-drug substitution as individuals cross the threshold; they find small, but significant, price elasticities that decline in drug cost, and full myopia ($\beta = 0$).

Our strategy builds on this literature to estimate elasticities with respect to variation in prices at multiple points in the budget set, for a broad range of the overall spending and age distributions. The method accommodates nonlinear price responses of the sort encountered in the energy economics literature, in which lack of information or understanding about prices leads consumers in nonlinear contracts (such as insurance contracts and electricity contracts) to use rules of thumb beyond discounting in determining their consumption. Liebman and Zeckhauser (2004) note that individuals may respond to nonlinear price schedules by “ironing” (responding to average or local average prices), or by “spotlighting” (responding to immediate “spot” prices rather than to future marginal prices). In the electricity setting, Ito (2014) finds that consumers respond to average price rather than marginal or expected

marginal price. While the empirical distribution of Medicare Part D spending shows clear evidence of bunching at kink points, which is not observed in Ito’s setting, it remains that there are several similarities between the energy setting and consumption under a nonlinear insurance contract: consumers may find it difficult to track their consumption to date and they have to form expectations regarding future consumption in order to determine their marginal prices. For this reason, we estimate a reduced form model that can accommodate average price, marginal price, and current-future price responses.

Finally, we consider decision-making in a complex setting by an elderly population. Issues considered in behavioral economics, such as myopia and salience, may be particularly acute among the elderly given that the potential for cognitive failures rises at older ages. A recent study by Agarwal, et al. (2009) shows that in ten different contexts, ranging from credit card interest payments to mortgages to small business loans, the elderly pay higher fees and face higher interest rates than middle-aged consumers.⁷ Several studies of these issues apply specifically to the context of Part D. Heiss, McFadden and Winter (2006); Kling, Mullainathan, Shafir, Vermeulen and Wrobel (2008); Abaluck and Gruber (2011); and Ketcham, et al. (2012) each study plan choice under Medicare Part D and find striking evidence in a variety of settings that elders do not make cost-minimizing choices of Part D plans (there is some disagreement regarding whether choices improve over time).

Price “salience” may be particularly important in this setting. Empirical studies of taxation and insurance have shown that consumers in complex or nontransparent environments can respond suboptimally to prices. For example, Feldman, Katuscak, and Kawano (2015) find that consumers responded to the removal of the lump sum Child Tax Credit by reducing their labor income, even though the removal did not impact their net wages; Abaluck and Gruber (2011) find that consumers are more responsive to Medicare Part D plan generosity variation that occurs through the premium than generosity that occurs through drug out-of-

⁷See also Salthouse (2004), which shows clear evidence that the performance on a series of memory and analytic tasks declines sharply after age 60; and Fratiglioni, Ronchi, and Torres (1999) for evidence on the relationship between the onset of dementia and age.

pocket costs. Crucially, such behavior has been found to depend on consumers’ knowledge and how information is presented. In Loewenstein, et al. (2013), only 14% of respondents were able to answer correctly four multiple choice questions about the four basic components of traditional health insurance design: deductibles, copays, coinsurance and maximum out-of-pocket costs; further, understanding of certain characteristics was worse than others (e.g., coinsurances were less well understood than copays). Similarly, Handel and Kolstad (2015) find evidence that individuals with employer-based health insurance misunderstand key features of plan generosity, with substantial financial consequences when plans are chosen. Chetty, Looney, and Kroft (2009) show that consumers are more responsive to commodity taxes when they are included in posted prices. In the Part D setting, we draw on this literature and allow for consumers to respond to more salient price changes over and above how those changes impact out-of-pocket prices faced by individuals. Our results suggests that perhaps the same features that lead elders to make errors in financial choices or in choosing the appropriate Medicare Part D plan lead them also to deviate from rational, forward-looking behavior in responding to cost-sharing features.

3 Identification Strategy and Data

3.1 Identification

The ideal variation to identify the impact of budget sets on consumption would include independent variation in each segment of the budget set and random assignment of individuals across plans. Unfortunately for our study as well as all others using Medicare Part D data, prices are endogenous for several reasons. First, prices result partly from the consumers’ decision of which plans to choose in light of their expected drug needs – even in the presence of the potential cognitive failures described above, sicker enrollees may choose more generous coverage. Second, prices chosen by pharmaceutical companies rise and fall in response to changes in consumer demand. Third, the non-linear budget set means that marginal prices

mechanically depend on consumption – if the price increases after the donut hole threshold, we expect to see a mechanical positive relationship between out-of-pocket (OOP) price and consumption, since sicker individuals are more likely to end up in the donut hole.

To deal with the first and second issues, we instrument for prices using variation generated by changes in Part D plan cost sharing rules, following the typical approach of health economics demand studies such as Chandra, et al. (2010). This approach will be biased to the extent that individuals switch plans due to cost sharing changes. We address this by creating the instrument using the change in cost-sharing in the initial plan of the enrollee, regardless of their ultimate plan choice. Under the assumptions (A) that individuals did not choose their initial plans based on anticipated changes in plan cost sharing the next year, and (B) that plans do not alter plan generosity in response to unobservable characteristics of enrollees that predict differential trends in consumption, cost sharing changes in the initial plan are exogenous to plan choice and to enrollee characteristics.⁸ We consider these assumptions to be quite reasonable in this setting. Regarding (A), Medicare Part D consumers are very insensitive to out of pocket costs even in the current year (Abaluck and Gruber (2011); Ho, Hogan and Scott-Morton (2015)) and so a fortiori are unlikely to choose plans based on anticipated future changes in coverage. Regarding (B), we control for a rich set of individual demographic and spending characteristics (including data on spending patterns in Medicare Parts A and B, which plans would not be privy to) in order to enable an apples-to-apples comparison of spending trends for individuals in plans that do and do not alter their coverage characteristics. While plans with different sets of enrollees may differentially change the coverage they offer, and while they may even do so in a way that is correlated with changes in consumption, our identifying assumption is only violated if they do so in response

⁸The very low rate of plan switching from year to year (roughly 10%) implies that this restriction to initial plan does not have much impact; indeed, as discussed in Section 5.3, our results are very similar if we exclude switchers from the analysis. See Ho, Hogan, and Scott Morton (2015) for a discussion of switching behavior. A more subtle problem is that the monotonicity constraint might be violated if, for example, “switchers” respond to coverage generosity decreases in their initially chosen plan by switching to a plan even more generous than their initial choice. We cannot test the monotonicity assumption directly, but as noted above we find that our results are not driven primarily by the behavior of active switchers – the coefficient estimates are similar between the full sample and the sample of non-switchers only.

to aggregate characteristics of enrollees not included in our rich set of controls for patient claims and expenditures.

The standard approaches to the third problem – non-linear budget sets – are either to estimate a nonlinear structural model assuming a particular model of optimization behavior as in Hausman (1985) or, more recently, Kowalski (2015) and Einav, Finkelstein, and Schrimpf (2015); or to estimate a nonparametric model with higher order terms for each segment and threshold of the budget set as in Blomquist and Newey (2002). In our analyses, we employ a simplified version of Blomquist and Newey by considering the linear response to budget set segments and by limiting our sample to individuals who are extremely likely to end the year well in the interior of a budget set segment. Robustness checks using higher order polynomial terms to isolate individual phase price responses in a manner analogous to Blomquist-Newey show similar patterns. We then relate the price responses from the flexible estimation specification to several candidate models of consumption behavior and examine which models perform better in terms of predicting in- and out-of-sample consumption.

The structural analysis serves several purposes. First, it allows us to explore counterfactuals – the impact of completely filling in the donut hole will depend on whether beneficiaries respond to nominal features of plans (i.e. “does this plan offer donut hole coverage?”) or the actual price they would face conditional on being in the donut hole. Second, the clinical implications would for example be quite different if the impact of the donut hole arises primarily from beneficiaries stopping consumption the moment they enter the donut hole vs. consuming more conservatively throughout the year if they expect to hit the donut hole. Third, understanding the underlying mechanisms provides a test of myopia and salience in a setting of great economic interest.

Our identification approach also has several nice features. First, we have variation in prices in both the initial coverage phase and donut hole. Over 90% of Medicare Part D enrollees end the year in one of these two phases, so this allows us to identify a marginal price response for enrollees over a wide range of total drug expenditures, rather than focusing on

behavior around the convex kink in the budget set at the donut hole for price variation. Second, variation in both “current” and “future” price as enrollees spend more over the course of the year allows us to estimate “current” and “future” price elasticities separately in our dynamic analyses and thus to determine whether consumers are primarily forward-looking or primarily myopic. Aron-Dine, et al. (2015) separately identify myopia and static price elasticities by comparing their future price elasticity estimates with price elasticity estimates calibrated from the RAND experiment; our variation allows us to make this comparison without relying on any external calibrations. Finally, relying on price changes within contracts for existing enrollees allows us to simulate price responses to policies that are most relevant for Medicare Part D enrollees: those in which existing plans’ generosity is altered but consumers are not forced to re-enroll. In contrast, relying on price responses for new enrollees would run the risk of over-estimating consumers’ price sensitivity to changes in plan design once they are enrolled.

3.2 Data and Variable Construction

We analyze data from a 20% sample of Part D enrollees from 2006 through 2009. The claims data contain information on drugs consumed, date of claim, quantity consumed (measured in days’ supply purchased on the claim – this is our outcome variable in all specifications), total retail price, and out-of-pocket price for each individual claim. The beneficiary data contain demographic variables and enrollment details. The plan and tier files contain detailed information on drug coverage in each coverage phase as well as nonlinear threshold information. In order to control as richly as possible for heterogeneity across individuals, we merge the Part D data with data on utilization in Medicare Parts A and B as well.

For our main analyses, we exclude individuals under 65, individuals who ever received low-income subsidies (and who thus were not subject to the majority of cost-sharing variation), and individuals who were enrolled in employer-sponsored plans. We focus on enrollees in standalone PDPs only. We analyze data for individuals enrolled in their Part D plan for the

full year in each year pair of analysis and who had at least one claim in each year. There are 451,632 sample enrollees in 2006-7; 1,126,682 sample enrollees in 2007-8; and 1,129,200 enrollees in 2008-9. 2006-7 included 1,372 plans, while 2007-8 and 2008-9 each included over 1,700 plans.⁹ Summary statistics on sample plans and enrollees are in Table 1. The majority of sample enrollees are white and female, with a mean age 75. Between the first and second year of each year pair, a small proportion (9-11%) of enrollees switched plans.

Table 1: Summary Statistics

	2006-7 Sample		2007-8 Sample		2008-9 Sample	
	Mean Sample Characteristics					
Num. Beneficiaries	451,632		1,126,682		1,129,200	
Num. Plans	1,372		1,720		1,722	
% White	95.93%		94.94%		94.92%	
% Black	2.35%		3.05%		3.05%	
% Female	62.28%		62.94%		62.53%	
Age	74.52		75.30		75.29	
% Switchers	11.06%		9.65%		9.98%	
	Standard Plan Features					
	Year 1	Year 2	Year 1	Year 2	Year 1	Year 2
Deductible	250	265	265	275	275	295
Donut Threshold	2,250	2,400	2,400	2,510	2,510	2,700
	Sample Plan Characteristics					
	Year 1	Year 2	Year 1	Year 2	Year 1	Year 2
% No Deductible	73.41%	78.65%	77.60%	79.51%	79.43%	79.06%
% Standard Deductible	24.04%	19.30%	20.48%	19.65%	19.58%	17.66%
% Full Donut Coverage	6.09%	2.18%	1.43%	0.00%	0.00%	0.00%
% Any Donut Coverage	13.74%	19.82%	16.35%	14.18%	13.83%	13.05%
	Sample Consumption					
	Year 1	Year 2	Year 1	Year 2	Year 1	Year 2
Mean Q	1,233	1,409	1,220	1,327	1,274	1,344
(Std. Dev.)	827	867	836	863	858	865
Mean Exp.	\$2,108	\$2,355	\$1,981	\$2,078	\$1,997	\$2,121
(Std. Dev.)	\$2,187	\$2,688	\$2,380	\$2,729	\$2,626	\$3,108

Notes: Table displays enrollee and plan summary statistics for full 20% random sample of unsubsidized, elderly Part D enrollees enrolled continuously in standalone PDPs throughout both years of each year pair (Medicare Advantage and employer-sponsored enrollees excluded). In addition, we exclude from our sample enrollees with zero claims in either year of the year pair.

The standard plan thresholds moved in each year of the program; the standard deductible increased from \$250 to \$295 between 2006 and 2009, and the standard donut threshold increased from \$2,250 to \$2,700. However, as noted above, the majority of enrollees were not enrolled in standard Part D plans. Between 18 and 24% of enrollees were in plans with the standard deductible, but 70-80% of enrollees were in plans with no deductible,

⁹The sample in 2006-7 is smaller due to our focus on individuals enrolled for the full year of each year pair. 2006, the first year of our sample and of the Part D program, had an extended enrollment period through May. Our results are not sensitive to the inclusion of individuals enrolling later in 2006.

and a small fringe of enrollees were in plans with positive, but nonstandard, deductibles. Furthermore, some enrollees had coverage in the donut hole; between 2006 and 2009, the proportion of sample enrollees with any donut hole coverage ranged from 13 to 20%.

Sample enrollees purchased 1,200 to 1,400 days' supply of prescription drugs per year on average, for a total expenditure (out-of-pocket plus plan expenditure) of about \$2,000 to \$2,400 per year. Note that, due to the extended enrollment period in the first year of the program, individuals enrolled throughout the entirety of 2006 had higher consumption than the average enrollee in later years, as expected if sicker enrollees signed up earlier in 2006.

Our analyses require a single actual price and price instrument for each enrollee, for each coverage phase, for each year of each sample year pair. We construct actual prices and price instruments in each coverage phase using plan coverage information at the coverage phase-drug (NDC) level, and aggregate those phases using enrollee-specific quantity weights based on days supply of drugs consumed.¹⁰ For year pair (year 1, year 2), the actual price in phase c of year y is the weighted average price the individual would face in phase c given the year y plan's year y cost-sharing rules; weights use the individual's year y consumption (days supply) across all sample drugs observed. That is, the price P_{icy} for individual i enrolled in plan p in coverage phase c , for year y of year pair year 1-year 2, is defined as

$$P_{icy} = \sum_{d \in D_{i,cs}} CS_{dcy,p} * RP_{dy,p} * w_{idy} + \sum_{d \in D_{i,cp}} CP_{dcy,p} * w_{idy}$$

where $CS_{dcy,p}$ and $CP_{dcy,p}$ are coverage phase-specific coinsurance rates and copays, respectively, for drug d in plan p , $RP_{dy,p}$ is retail price for drug d in plan p , and the consumption weight for drug d is calculated as $w_{idy} = q_{id,y} / \sum_{d \in D_{i,y}} q_{id,y}$ using observed quantity consumed $q_{id,y}$ for each individual-drug-year combination. $D_{i,y}$ is the set of all drugs consumed by individual i in year y . Prices are for 30-day supplies of drugs. The retail price for a given plan-drug-year combination is calculated as the average retail price (total expenditure per

¹⁰See Appendix A for a detailed example.

30-day supply) across all observations of that plan-drug in the claims data for that year.¹¹

The price *instrument* in phase c of year y is the weighted average price the individual would face in phase c given the year 1 plan's year y cost-sharing rules; weights use the individual's year 1 consumption. For coinsurances, we apply the coinsurance rate to the retail price appropriate for the given plan-drug combination in year 1. That is, the IV price P_{icy}^{IV} is defined as

$$P_{icy}^{IV} = \sum_{d \in D_{i,cs}} CS_{dcy,p(year1)} * RP_{d,year1,pooled} * w_{id,year1} + \sum_{d \in D_{i,cp}} CP_{dcy,p(year1)} * w_{id,year1}.$$

Variation in the instrumental variable prices in the first year of each year pair are shown in the top panel of Table 2. Note that these prices are per 30 days supply; the average beneficiary in our sample purchases 1,250 – 1,350 days supply of drugs per year, so the impact of even a small change in prices per 30 days supply is substantial. As in the standard plan, the average price decreases, then increases, then decreases again as one moves from the deductible to the initial coverage range (ICR), from the ICR to the donut hole, and from the donut hole to the catastrophic phase. Average differences between phases are not as large as they would be in the standard plan (which has 100% coinsurance in the deductible and donut, 25% coinsurance in the ICR), because many enrollees have no deductible (in which case the “deductible” price in the Table is effectively the ICR price), and some enrollees have coverage in the donut hole.

Year-to-year changes in individuals' IV prices are shown in the bottom panel of Table 2; price changes are shown broken out by coverage phase and year pair. Note that, because plan choice, consumption weights, and retail prices are held fixed at year 1 values, conditional on those year 1 variables, price differences are a function only of year 1 plan cost-sharing changes between year 1 and year 2. For the sake of brevity, price differences are shown only for the

¹¹When possible, claims for 30-day supplies only are used to calculate average retail prices. When 30-day supply claims are not observed for particular plan-drug combinations, retail price per 30-day supply is imputed by scaling average prices per one-day supplies observed in claims for all other quantities.

ICR and donut hole.¹² There is considerable variation in year-to-year ICR price changes across sample individuals; the sample of individuals experiencing price changes in the donut hole is smaller since the IV price change is zero by construction for those individuals with no gap coverage in both years. Price elasticities are identified by comparing consumption trends across similar individuals experiencing different price trends; thus all individuals, including those with and without some gap coverage, are used to identify model parameters.

Table 2: Sample Price Instrument Variation

	2006-7 Sample		2008-9 Sample		2008-9 Sample	
	Mean	(SD)	Mean	(SD)	Mean	(SD)
<i>Year 1 IV Prices</i>						
Deductible	26.81	(31.43)	23.04	(34.65)	21.88	(44.22)
Initial Coverage Range	16.58	(14.73)	15.41	(15.76)	14.50	(28.47)
Donut Hole	50.80	(60.12)	49.62	(60.02)	49.19	(98.43)
Catastrophic Range	3.99	(2.98)	3.95	(2.89)	3.90	(4.83)
<i>Year 2 IV Price - Year 1 IV Price</i>						
Initial Coverage Range	1.01	(7.93)	1.32	(3.49)	3.19	(7.33)
Donut Hole	5.64	(36.77)	7.10	(20.04)	1.28	(16.54)

Notes: Price instruments generated using plan-drug-coverage phase-specific cost-sharing parameters (copays and coinsurances) and individual enrollee-specific consumption weights on drugs. For prices specified as coinsurances, average retail price for each plan-drug combination used as the basis to which plan-drug-phase coinsurances are applied. Enrollee-specific consumption weights based on days supply used to generate a weighted average price for each individual in each coverage phase. For second year of each year pair, consumption weights, retail prices, and copays/coinsurances from first year consumption and plan enrollment imposed to isolate price effect of changes in cost-sharing characteristics holding consumption and enrollment behavior fixed. Donut price change shown only for plans with any gap coverage in either year (16% of sample overall; price instrument difference is mechanically zero for other plans).

In looking across the three year pairs of our analysis sample, we note several patterns of interest. First, we see in the bottom panel that the ICR and donut price changes are slightly positive on average, though the standard deviation indicates that there are many substantial positive and negative price changes in the sample. That is, within a given contract, the average enrollee experiences diminishing plan generosity for their fixed bundle of drugs between years 1 and 2. On the other hand, a comparison of prices across years in

¹²Individuals with deductibles in both years of the year pair have no effective IV variation in the deductible because retail prices are held fixed across years. Similarly, catastrophic price variation is based only on retail prices and cost-sharing minimums, the former of which are held fixed between years in the IV and the latter of which vary across years, but not across plans. Thus, the majority of our variation within coverage phase comes from the ICR and donut hole cost-sharing changes. In some analyses below, we also analyze responses to variation coming from changes in the location of deductible thresholds between years.

the top panel of Table 2 indicates that average prices across existing *and* new plans actually decrease between 2006 and 2008. Thus, changes in generosity over time impact existing enrollees less favorably than new enrollees (a similar pattern applies to premiums for new vs. existing enrollees, as documented in Ericson (2014) – this phenomenon is termed invest-then-harvest behavior in the context of premium-setting). Second, the variation in the IV donut price change is falling over time, which is consistent with the donut hole data in Table 1 – fewer enrollees have donut hole coverage in 2009 than in 2006, and no sample enrollees have full donut coverage in 2009, as opposed to 6% in 2006. Finally, we note that the magnitude of the price variation generated by our instrument is substantial. Comparing, e.g., the 5th and 95th percentile ICR price change in 2006, a difference of \$12.63 per 30 days supply would translate to more than a \$500 difference in OOP costs given mean year 1 consumption.

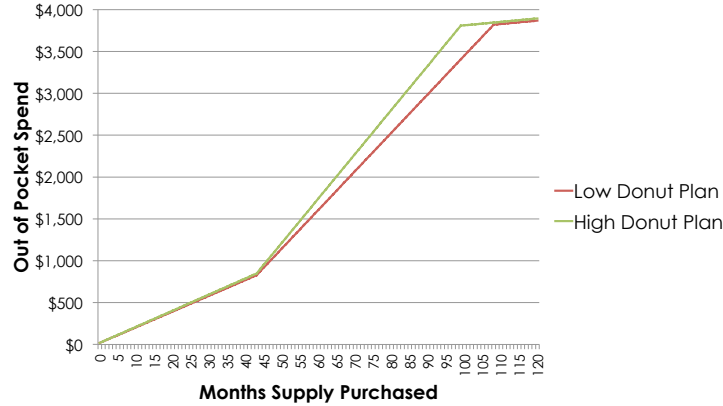
3.3 Illustrative Example – Donut Hole Coverage

As an illustrative example of our identification strategy, we start by considering how individuals respond to changes in donut hole coverage, all else equal. To isolate the consumption response to the donut hole coverage change only, we take all the plans in 2006-7 and match based on having the exact same coverage thresholds in both 2006 and 2007 and the same donut hole coverage in 2006, and we impose that matches have similar prices in year 1 and prior to the donut hole in year 2. Within those matched plans, we select matched pairs that have *different prices in the donut hole* in year 2.¹³ Taking unique plan matches only, we obtain a large sample of 97,330 individuals in matched plans, where plans that either dropped generic coverage between 2006 and 2007 or added it were matched to plans which did not change their donut hole coverage between years.¹⁴

¹³In the results shown, we have imposed that: the difference between plans' donut prices in 2007 exceeds \$3; that the difference between plans' deductible prices in 2007 is less than \$1; and that the difference between plans' 2006 weighted average prices (using pooled sample consumption weights based on days supply) and the difference between plans' 2007 ICR prices are each less than one-third the 2007 donut price difference between plans. Results are not sensitive to changes in these matching parameters, though more restrictive matches have larger standard errors.

¹⁴Many plans have multiple matches based on the criteria above. To obtain unique matches, we iterate over matches based on size – we find the best match (the match with the most similar 2006 and pre-donut

Figure 1: Matched Plans – OOP Cost Comparison (2007)



Notes: Average cumulative out-of-pocket (OOP) spending as a function of cumulative months-supply purchased based on prices in 188 matched plans. Plans are matched based on similarity in both 2006 coverage characteristics and 2007 pre-donut coverage characteristics and dissimilarity in 2007 donut coverage. Details in text. "Low Donut" plans are those with lower prices (generic coverage) in the donut hole in 2007 (average OOP price \$46); "High Donut" plans are those with higher prices (no coverage) in the donut hole in 2007 (average OOP price \$53).

Figure 1 shows a comparison of cumulative out-of-pocket cost as a function of total days supply purchased for “high donut” vs. “low donut” matched plans. The plans within each pair have similar ICR cost-sharing (\$19 per 30-day supply on average); however, across all plan pairs, the average out-of-pocket cost in the donut hole for a 30-day supply of drugs is \$53 in the “high donut” plan vs. \$46 in the “low donut” plan. Accordingly, we see that cumulative out-of-pocket cost is equivalent in the high- and low-donut plans until the donut threshold is reached, at which point the high-donut plan’s out-of-pocket cost increases more steeply.¹⁵ Note that the “low donut” plans are still less generous in the donut hole than in the ICR (there is still a “kink” at the donut threshold) because they at most only cover generic drugs in the donut hole.

Using this sample, we ran a differences-in-differences regression of the year-to-year annual or quarterly change in quantity purchased on plan pair fixed effects and a dummy for being

2007 prices and least similar 2007 donut prices) for the largest matched plan (by enrollment), remove both plans from the sample, and iterate on this procedure until each match is unique. In the final sample, we have 94 matched pairs with 188 plans total.

¹⁵The plans have the same catastrophic threshold, but the catastrophic threshold is reached more quickly in high-donut price plans, because crossing the catastrophic threshold is defined based on cumulative out-of-pocket costs rather than on total expenditure. Both the deductible and donut thresholds in Part D plans are reached based on expenditure. None of the matched plans have deductibles.

in the “high donut hole price” plan in the second year (2007):

$$Q_{2it} - Q_{1it} = \alpha + \beta * d^{HighDonutChg} + \delta * X_{ij}$$

These regressions include controls for individual demographics and individual health care consumption patterns. The former include dummies for age, sex, race, and state. The latter include rich controls for total base year (2006) individual drug consumption and prices (polynomial controls for coverage phase-specific prices, and 100 quantiles each of days supply purchased, total drug expenditure, out-of-pocket (OOP) drug expenditure, and retail price per prescription) and total base year (2006) individual medical spending (dummies for nonzero spending overall and in several subcategories – office visits, inpatient emergency, inpatient non-emergency, outpatient, and other – as well as polynomial controls for medical spending overall and level controls for spending in each subcategory). We also include dummies for having any drug spending in each generic therapeutic class (GTC), which is a summary measure of target medical condition.¹⁶ We interact the GTC dummies with separate indicators for generic and branded drugs. These controls incorporate the possibility that underlying utilization trends may differ by type of illness and preferences over branded vs. generic drugs. Finally, we include dummies for each plan pair (noting that each plan in a plan pair has the same year 1 budget set and pre-donut year 2 budget set). Using this specification, we mimic the thought experiment from Section 3.1 – we find plans that are close to identical in year 1 and use rich controls so that identification of price responses comes from comparing similar individuals, some of whom experience price changes.

The results in Table 3 show several striking facts. First, enrollees’ consumption responds negatively to lower donut hole coverage. Second, the response is large (-0.085) and significant

¹⁶Each unique drug (NDC) is classified by the company First Data Bank as falling under a particular generic therapeutic class based on the medical condition it treats. There are forty such classes, the most popular of which are “Cardiovascular,” “Autonomic Drugs,” “Cardiac Drugs,” and “Diuretics” among our sample enrollees. The GTC by generic dummies account for the potential that, for example, individuals who tend to take cardiac drugs may exhibit different utilization trends than those who take anti arthritics, even absent differential price changes between sample years.

at the end of the year, when people are more likely to have entered the donut hole, while it is undetectable in the first quarter of the year.¹⁷ Myopia is one possible explanation for this pattern – in Section 5, we present more direct evidence that it plays a substantial role.

Table 3: Differences-in-Differences Results, Donut Hole Coverage Dummy, Matched Plan Enrollees

	Obs	Coef	SE
Full Year	97,330	-0.025	0.016
Q1	78,867	0.002	0.018
Q2	78,867	-0.031	0.015 *
Q3	78,867	-0.048	0.015 **
Q4	78,867	-0.085	0.016 **

Notes and Sources: Results of full year and quarterly regressions of 2006-7 consumption change on a dummy for being in a "high year 2 donut" plan. January enrollees only. In full year runs, individuals with zero quantity in either year dropped. In quarterly runs, individuals with zero consumption in any quarter of either year dropped. Regression controls in text. Superscript "***" indicates significance at the 1% level; superscript "**" indicates significance at the 5% level.

This illustrative example involves a stark change in coverage in one region of the Part D plan budget set. In order to translate the results into a model of consumption behavior, we must consider how observed responses at a point in time relate to current prices and prices not yet encountered in the budget set, and how they vary with actual out-of-pocket costs as opposed to categorical changes in coverage that are more salient to consumers.

4 Motivating a Flexible Empirical Model

In this Section, we outline a baseline model of optimal consumption with nonlinear budget sets and describe several ways in which consumption may deviate from this model in practice. We then present a flexible empirical model that can accommodate these deviations.

We begin with a simple model, in which forward-looking, rational individuals optimize their consumption of a single prescription drug over the course of the year with no un-

¹⁷Note that we do not necessarily expect the full year coefficients to be a mechanical average of the quarterly coefficients, because the specification is run in logs and because the quarterly analyses are restricted to individuals with positive claims in each quarter and thus to a higher-utilization baseline enrollee. Qualitative patterns are insensitive to this restriction.

certainty. Suppose individuals maximize their utility $u(\cdot)$ by choosing consumption q_t for each $t = 1, \dots, T$. Suppose also that they face a nonlinear budget set with the following out-of-pocket expenditure function over total quantity $Q = \sum_{t=1}^T q_t$ purchased in the year:

$$E^{OOP}(Q) = \begin{cases} p_1 * Q & \text{if } Q \leq \frac{\bar{X}}{R} \\ p_1 * \frac{\bar{X}}{R} + p_2 * \left(Q - \frac{\bar{X}}{R}\right) & \text{if } Q > \frac{\bar{X}}{R} \end{cases}$$

where out-of-pocket prices p_1 and p_2 are indexed by coverage phase (p_1 being the first coverage phase encountered, p_2 being the second), there is a single coverage threshold at \bar{X} , and the drug's retail price (total price paid by the plan *plus* enrollee) is R . As in many nonlinear budget sets, the effective price paid by consumers changes as a function of total spending: for the first $\frac{\bar{X}}{R}$ units, the unit price is p_1 ; for the remaining $\left(Q - \frac{\bar{X}}{R}\right)$ units, the unit price is p_2 . For example, \bar{X} could represent a convex budget set kink with out-of-pocket prices $p_1 < p_2$ as would occur at the Medicare Part D donut threshold.¹⁸

Given perfectly forward-looking behavior, no uncertainty or discounting, and $u(\cdot)$ constant across t , the dynamic problem collapses to the static choice of prescription drug quantity Q for the full year. The solution for a general quasiconcave, continuously differentiable utility function over prescription drugs (assuming linear utility for the numeraire) will be:

$$Q^* = \begin{cases} \tilde{u}'^{-1}(p_1) & \text{if } \tilde{u}'\left(\frac{\bar{X}}{R}\right) \leq p_1 \\ \frac{\bar{X}}{R} & \text{if } p_1 < \tilde{u}'\left(\frac{\bar{X}}{R}\right) \leq p_2 \\ \tilde{u}'^{-1}(p_2) & \text{if } \tilde{u}'\left(\frac{\bar{X}}{R}\right) > p_2 \end{cases} \quad (1)$$

where $\tilde{u}(Q) = \sum_{t=1}^T u(\frac{Q}{T})$. In the optimum, individuals consume either as they would under a linear price schedule with p_1 or p_2 or they “bunch” at the coverage threshold \bar{X} .¹⁹ This

¹⁸In a 2006 Part D plan with standard cost-sharing and no deductible, for a single \$100 drug, we would have $p_1 = \$25$, $p_2 = R = \$100$, and $\bar{X} = \$2,250$. We focus on response to ICR and donut hole prices in this project, as very few individuals end the year in the deductible or catastrophic phase in Medicare Part D. However, the model easily generalizes to accommodate more coverage phases.

¹⁹If we expand the model to allow for kinks at the deductible and catastrophic thresholds, the model optimum is similar, but there is no “bunching” at deductible or catastrophic thresholds. At thresholds

makes intuitive sense: those consumers whose marginal utility of consumption at \bar{X}/R is less than p_1 prefer to consume below the threshold \bar{X}/R at linear price p_1 , and would continue to do so for any p_2 with $p_2 \geq p_1$; their consumption is thus unaffected by the nonlinearity of the budget set. Similarly, those consumers whose marginal utility of consumption at \bar{X}/R exceeds p_2 prefer to consume past the threshold under linear price p_2 , and would continue to consume beyond the threshold if prices for marginal *or* inframarginal units drop (the additional savings on inframarginal units are akin to a transfer given the assumed quasilinear utility function). Those consumers who would be willing to pay p_1 for an additional unit of consumption at the kink but who are not willing to pay the post-kink price p_2 for that unit will bunch at the donut hole by consuming exactly \bar{X}/R for the year.

In sum, in the somewhat restrictive model described above where consumers have perfect knowledge regarding (and are perfectly attentive to) their price schedule and prescription drug needs throughout the year, we predict that total consumption for many price-parameter combinations will be a function of marginal (end of year) price as in a linear budget set model. Next, we relax the model to allow for dynamic decision-making due to myopia and account for the role of uncertainty.

4.1 Dynamic Consumption Responses

Now, we modify the model slightly and assume that individuals make their prescription drug utilization choices in each of T periods. Letting V_t denote the value function from period t forward given previous expenditure X_t , and allowing for hyperbolic discounting, we arrive

where the out-of-pocket price decreases, the individual would prefer to “jump” past the threshold – if an individual’s marginal utility for an additional unit of consumption exceeds the pre-threshold price p_1 , then it also exceeds $p_2 < p_1$.

at the following program:

$$\begin{aligned}
V_T(X_T) &= W_T(X_T) = \max_q u(q) - E^{OOP}(X_T, q) \\
V_t(X_t) &= \max_q u(q) - E^{OOP}(X_t, q) + \beta W_{t+1}(X_t + q * R) \quad \forall t < T \\
W_t(X_t) &= \max_q u(q) - E^{OOP}(X_t, q) + W_{t+1}(X_t + q * R).
\end{aligned}$$

Note that, in each $t < T$, total utility in all future periods W_t is discounted by β – we allow for hyperbolic discounting as in a $\beta - \delta$ discounting model, but impose that $\delta = 1$. This is a reasonable assumption, as the literature on financial decision-making generally finds long-run annualized discount factors to be very close to 1; see Laibson, et al. (2015) for a review.²⁰ Consumption in the final period T will be the solution to the static optimization problem as in equation (1) for a nonlinear budget set with a kink at $\tilde{X} = \max\{\bar{X} - X_T, 0\}$. Consumption in $t = 1, \dots, T - 1$ will instead maximize current-period utility $u(q)$ less out-of-pocket expenditure E^{OOP} plus (discounted) future utility; current consumption impacts future value W_{t+1} via the nonlinear budget set.

If we further assume that $u(\cdot)$ is quadratic:

$$u(q) = \tilde{\alpha}_0 + \tilde{\alpha}_1 * q + \tilde{\alpha}_2 * q^2$$

then the solution will be such that, for all individuals ending both the current period and the year in the interior of a coverage phase, consumption in each period is given by:

$$q_{it} = \alpha_{1t} + \alpha_{2t}(\beta MP_{i,y} + (1 - \beta)CP_{it,y}) \quad (2)$$

with $\alpha_{1t} = -\frac{\tilde{\alpha}_{1t}}{2\tilde{\alpha}_{2t}}$ and $\alpha_{2t} = \frac{1}{2\tilde{\alpha}_{2t}}$. The terms $MP_{i,y}$ and $CP_{it,y}$ denote the marginal end of year price and the current price in period t , respectively. Consumption in each such period

²⁰We use hyperbolic discounting to capture “myopia” in the form of present-biased price responses, but we note that such responses could be the result of multiple models of decision-making, including one in which consumers are not attentive to the full schedule of prices that they face (as opposed to being perfectly attentive, but valuing future returns lower than current returns).

will be a function of a weighted average of current price and future price, where the weight on future price equals the discount factor β . Hence, the coefficients on current price and marginal price will equal the overall price sensitivity term α_{2t} multiplied by the weights $1 - \beta$ and β , respectively, due to discounting. This parallels the current-future price specification in Aron-Dine, et al. (2015).

This specification can easily be relaxed to allow for more general classes of utility function and uncertainty regarding prescription drug needs. For example, suppose utility in each period is instead $u_t(q) = u(q, A_t)$, where A_t denotes a (stochastic) set of parameters impacting the utility function in period t . In Appendix E, we derive and prove a simple theorem characterizing optimal consumption under uncertainty, in which consumption in periods prior to the final period T are optimized given the individual's *expectation* of hitting or crossing the budget set threshold \bar{X} in future periods. That is, consumption is determined as a function of current price and expectations regarding future prices rather than current price and actual marginal price.

4.2 Allowing for More Flexible Models of Consumption

In practice, current and future prices may be complicated objects for individuals to calculate, and consumers may further be confused about the role of future prices in their optimal consumption path. As noted above, research has found that in other settings with nonlinear pricing (such as electricity markets), consumers may respond to average prices rather than to expected marginal prices. In order to account for the fact that the current-future model derived above may impose the wrong model of consumer behavior, we consider a more flexible specification of the dynamic utilization model which subsumes the current/future model without imposing a specific functional form on how individuals respond to different segments of the budget set. Specifically, we regress quarterly quantity consumed on the initial coverage range (ICR) price and the donut hole price.²¹

²¹The price variation analyzed in this study is primarily in the initial coverage phase and donut hole, but some plans also have nonzero deductible thresholds in one year or both; in subsequent Sections, we address

Returning to the rational model in which consumers respond only to marginal end of year prices ($\beta = 1$) and know exactly what these prices will be (no uncertainty), the extent to which consumers respond to prices in each coverage range of their budget set will depend on the probability that each coverage range is marginal. By breaking the marginal price MP down into its component parts, we can write:

$$q_{it} = \alpha_1 + \alpha_2 MP_{i,y} = \alpha_1 + \alpha_2 \mathbb{1}(m_i = ICR) * P_{ICR} + \alpha_2 \mathbb{1}(m_i = Donut) * P_{Donut}$$

where $\mathbb{1}(m_i = C)$ is an indicator for the event that coverage phase C is the marginal coverage phase for individual i . Rewriting this as

$$q_{it} = \alpha_1 + \alpha_{ICR} * P_{ICR} + \alpha_{Donut} * P_{Donut}, \quad (3)$$

the coverage phase-specific coefficients will be such that $E(\alpha_{ICR}) = E(\alpha_2 * \mathbb{1}(m_i = ICR)) = \alpha_2 * \Pr(m_i = ICR)$ and similarly for the donut hole price coefficient.²²

In the more general current/future model with myopia, the specification as a function of coverage phase prices becomes:

$$\begin{aligned} q_{it} &= \alpha_1 + \alpha_2 * \beta * (\mathbb{1}(m_i = ICR) * P_{ICR} + \mathbb{1}(m_i = Donut) * P_{Donut}) \\ &+ \alpha_2 * (1 - \beta) * (\mathbb{1}(c_i = ICR) * P_{ICR} + \mathbb{1}(c_i = Donut) * P_{Donut}) \\ &= \alpha_1 + \alpha_{ICR} * P_{ICR} + \alpha_{Donut} * P_{Donut}, \end{aligned}$$

where $\mathbb{1}(c_i = C)$ is an indicator for the event that C is the current coverage range for individual i . This implies that the coefficients on the ICR and donut respectively will be

$$E(\alpha_{ICR}) = \alpha_2 * (\beta * \Pr(m_i = ICR) + (1 - \beta) * \Pr(c_i = ICR)) \text{ and } E(\alpha_{donut}) = \alpha_2 * (\beta * \Pr(m_i =$$

this variation as well.

²²We are implicitly assuming here that the LATE estimated given our instrument will equal the ATE. This assumption is equivalent to asking whether compliers – consumers for whom the price change of their year t-1 plan impacts prices today – have systematically different marginal prices than non-compliers. If they do, the above probabilities will be the probabilities among compliers rather than among the whole sample.

$Donut) + (1 - \beta) * P(c_i = Donut))$. In other words, this reduced form model of consumption responding to ICR and donut prices subsumes a model of consumption responding to current and future price; however, we do not *require* that the consumption response to coverage phase prices be scaled by current and future probability weights under a specific model of expectations, and can accommodate behavior such as “ironing,” where individuals respond to average price rather than a combination of current and marginal prices.

Our identification approach allows us to obtain causal reduced form price responses, which will be unbiased by endogeneity between consumption and price induced by the nonlinear budget set. However, depending on the specification of the utility function, any structural relationship between coverage range prices and consumption will be complicated by endogenous coverage phase transitions and may change in the presence of uncertainty. In the linear demand specification (quadratic utility), the above result holds even if there is uncertainty about which coverage range is current/marginal as long as price changes do not alter the probabilities of coverage phase transitions and as long as there is no bunching. With other utility functions, this result no longer holds exactly.²³ We contend with these issues in three ways. First, in our baseline analyses, we attempt to minimize uncertainty, as well as the potential for endogenous coverage phase transitions in response to prices, by focusing our analyses on individuals predicted to end the year well away from kink points. We also examine the robustness of this approach using multiple analyses in Appendix E. We show that the estimates are not sensitive to more conservative sample restrictions. We also demonstrate the limited impact of uncertainty and endogenous phase transitions by re-estimating a richer version of the preferred structural model with endogenous future prices and uncertainty. As discussed in Appendix E, the results with these modifications are unchanged.

²³For example, our estimates will be the same in the linear model whether beneficiaries A and B both believe they have a 50% chance of hitting the donut hole and respond half as much as they otherwise would or whether A believes he has a 100% chance of hitting the donut hole and B a 0% chance. In the log model, we have $\log(C_i) = \alpha_1 + \alpha_2 \log(E_i(P_i))$ where $E_i(P_i) = \sum_c P(\text{coverage range}_i = c) P_{ic}$ where P_{ic} is the price in each coverage range.

4.3 Allowing for Price Salience

The above setup supposes that consumers will respond to cost-sharing variables only insofar as they impact the out-of-pocket prices that enter the budget set. In practice, however, choosing optimal consumption in a Part D plan requires not only dynamic optimization with a nonlinear budget set, but also a calculation of within-phase prices given the particular prescription drugs each enrollee takes (or may take). In the case of cost-sharing specified as copays, the enrollee must consult the formulary and plan benefit description for each drug and aggregate; in the case of coinsurances (as are generally used to determine cost-sharing at least in the deductible and donut hole), the enrollee must also know each drug’s retail price to determine his or her cost-sharing. Notably, although consumers may learn about retail prices and/or cost-sharing for the current coverage phase when they purchase a drug at the pharmacy, they may yet have residual uncertainty about prices in coverage phases they could encounter in the future.

As discussed in Section 2, a number of empirical studies demonstrate that consumers underreact to less salient indicators of price (e.g., insurance plan coverage or taxes) and overreact to more salient indicators of price.²⁴ To the extent that some portion of the calculation exercise described above makes the phase-specific average price for an individual enrollee’s bundle of drugs *less salient* than other indicators of plan coverage generosity, the current-future price specification that results from the above dynamic model (or alternative dynamic models with different phase weighting) may fail to adequately account for the full scope of individuals’ responses.

Some plan characteristics that may specifically be more salient are the presence or absence of a deductible, and the presence or lack of donut hole coverage, each being particularly visible in plan benefit materials and tools such as the Medicare Plan Finder on CMS’s website.

²⁴Salience effects contemplated in this Section may be a function of both limited awareness of plan characteristics and incomplete understanding of those characteristics’ implications for prices. Limited awareness has been found to be important in the health insurance setting; see, e.g., Loewenstein, et al. (2013), Handel and Kolstad (2015). Incomplete translation of salient characteristics into prices has a substantial impact in the tax context in Feldman, Katuscak, and Kawano (2015).

Regarding the deductible, we expect consumers to respond to the presence or absence of coverage either because they are myopic²⁵ and/or because the deductible is a particularly salient coverage benefit; in our setting, we cannot distinguish these two mechanisms because there is no within-phase price variation conditional on deductible coverage.

However, we can examine whether individuals respond to categorical donut hole coverage conditional on out-of-pocket price in the donut hole. Some of the donut hole coverage variation encountered by our sample enrollees is of the sort of clear variation used to estimate price sensitivity in previous work, such as the RAND experiment or Chandra, et al. (2010); for example, whether the plan includes coverage of branded drugs in the donut hole, or has no donut hole coverage. Other donut hole coverage variation is more subtle and would be difficult for enrollees to translate into prices; for example, whether the plan includes coverage of “many preferred” branded drugs. As discussed in Appendix C, we may expect consumers to respond to nominally large changes in coverage (denoted “stark changes” below), over and above how those changes translate into out-of-pocket price in the donut hole.

There is ambiguity about how best to model consumers’ responses to more salient characteristics that are not entirely captured in out of pocket prices. In the interest of simplicity, we extend the above empirical specification in equation (3) to include indicators for whether the plan includes a deductible, as well as whether it includes generous, easy-to-understand (“stark”) donut hole coverage.

$$q_{it} = \alpha_1 + \alpha_{ICR} * P_{ICR} + \alpha_{Donut} * P_{Donut} + \theta_{Ded} * ded + \theta_{Stark} * stark$$

In the following Section, we estimate this flexible model with and without the salience terms included; we then use the reduced form results to examine performance of the current-future model (vs. alternative models of expectations and weighting) in this setting.

²⁵The deductible is not the marginal price for all but a small sliver of the Part D population.

5 Reduced Form Price Response Estimates

In this Section, we present the results of several reduced form empirical specifications aimed at examining the mechanisms described in the previous Section: overall sensitivity of Part D enrollees to prices, myopia regarding future price changes, and price salience. We first model dynamic (quarterly) consumption as a function of ICR and donut prices, as well as changes in categorical coverage (in the deductible and donut hole regions of the budget set). The results indicate that individuals do respond negatively and significantly to prices overall, and that they are more responsive to current prices than to prices they will encounter in future periods. We also demonstrate that they are responsive to categorical coverage variation over and above how that variation impacts out-of-pocket prices.

5.1 Sample: Individuals Away from Kink Points

The dynamic specification outlined above only holds for individuals ending both the current and marginal periods in the interior of a coverage phase – accordingly, this analysis focuses on individuals for whom the donut is highly likely to be marginal and individuals for whom the ICR is highly likely to be marginal (excluding individuals likely to be near the threshold between the two budget set segments based on previous year spending). Specifically, for the 2006-2007 analysis, we look only at individuals consuming less than or equal to \$1,500 in 2006 and, separately, at individuals consuming between \$3,000 and \$5,000 in 2006. We restrict the samples similarly for 2007 and 2008, adjusting the cutoffs in each year pair to account for secular trends in standard plan thresholds. The first group is a set of individuals who are almost certainly going to have the ICR price as their marginal price; the second group is a set of individuals almost certain to have the donut price as their marginal price.

In practice, 3.7% of individuals in the low-spending group cross the donut hole threshold in year 2 and 14.2% of individuals in the high-spending group do not. Thus, our sample restrictions do not completely eliminate switching behavior and may be impacted by un-

certainty. Further, the extent to which coverage phase switching responds endogenously to prices may bring an additional source of bias. However, we can use these samples to examine behavior in the presence of far more limited marginal price uncertainty and scope for switching than we would expect in the full sample. In Appendix E, we limit the sample further to eliminate bias due to switching uncertainty and estimate a more general model with endogenous switching. Our results are nearly identical to the baseline results, implying the linearized model with sample restrictions relative to budget set kinks performs quite well.

5.2 Dynamic Price Responses

We empirically estimate the model above, with consumption as a function of the ICR and donut price, assuming power utility (motivating a log functional form). The log functional form is employed because the distribution of consumption is skewed positive and there is a heavy upper tail even within the restricted regression sample. This yields an equation of the form: $\log(q_{it}) = \alpha_1 + \alpha_{t,ICR}\log(P_{i,y,ICR}) + \alpha_{t,Donut}\log(P_{i,y,Donut})$ for individual i in period t of year y . For each quarter and each pair of years, we take differences across years, yielding:

$$\begin{aligned} \log(q_{i,t,y}) - \log(q_{i,t,y-1}) &= \alpha_{1t} + \alpha_{t,ICR}(\log(P_{i,y,ICR}) - \log(P_{i,y-1,ICR})) \\ &+ \alpha_{t,Donut}(\log(P_{i,y,Donut}) - \log(P_{i,y-1,Donut})) + u_{it} \end{aligned}$$

As described in Section 3.2, in this and each of the following regressions, we use an instrumental variables strategy based on plan choice inertia. For a given pair of years, cost-sharing characteristics used to generate the year 2 price instrument are the year 2 cost-sharing parameters (copays/coinsurances, coverage thresholds, etc.) of the year 1 chosen plan. In all regressions, we include all the controls from the paired analysis as well as rich controls for plan characteristics in order to generate the apples-to-apples comparison described in Section 3.1 – two individuals with identical plans and observables in 2006, one of whom experiences a change in cost-sharing in 2007. As this exercise does not explicitly pair plans

based on year 1 plan characteristics, we also control for dummies for year 1 deductible and donut hole coverage and thresholds, polynomials of year 1 prices in each budget segment, and 50 quantiles each of year 1 quantity (days supply), expenditure, out-of-pocket spending, and average retail price of drugs consumed for the average person in the year 1 plan.²⁶

Consider first the low-spending group. Results are in Table 4.²⁷ The proportion of individuals in their marginal coverage phase is slightly increasing over the course of the year (beginning at 83% in Q1, ending at 97% in Q4) as the small proportion of individuals with deductibles enter the ICR.²⁸ Considering each year pair individually, the ICR price response is either flat (2006-7) or slightly increasing in magnitude over the course of the year (2007-8 and 2008-9). On balance, the results for all years pooled show that the ICR response is fairly flat across quarters even though the proportion of individuals in the marginal coverage phase increases slightly – given that most individuals (83%) either do not have a deductible or exit the deductible in Q1, these results are limited in their usefulness for detecting myopic behavior. We do not observe a substantial response of low-spending individuals with respect to the donut hole price, as would be expected given even imperfectly forward-looking behavior – in some samples, we observe a small positive sign on the donut hole price. The magnitude of the static price elasticities suggested by these estimates is on the lower end of the elasticities found in the literature (-0.04 to -0.05).

Consider next the high-spending group. Results are in Table 5. Among high-spending individuals, there is a steep increase in the proportion of individuals in the marginal coverage phase (the donut hole) between quarters 1 and 4 (rising from 0.5% to 95% based on year 1 consumption, the latter number reflecting that some high-spending individuals are in plans with no donut hole). Concurrent with this increase, we observe also that the donut hole price response is quite steep over the course of the year in each year pair individually and in the

²⁶In the paired analysis in the previous section, these plan characteristics would be collinear with the plan pair fixed effects and “high donut” dummy.

²⁷In these and all subsequent results, instrumental variables estimates are shown; see Appendix B for OLS, reduced form, and IV results and accompanying discussion.

²⁸The remaining 3% do not exit the deductible in year 1.

Table 4: Results of Quarterly ICR and Donut Price Regressions – Low-Spending Group

Period	Price	% in ICR	All Years Pooled		2006-7		2007-8		2008-9	
			Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	83.0%	-0.051	0.006 **	-0.047	0.014 **	-0.022	0.010 *	-0.061	0.009 **
Q1	Donut		0.023	0.008 **	-0.008	0.023	0.022	0.011 *	0.018	0.010
Q2	ICR	92.1%	-0.041	0.007 **	-0.046	0.018 *	-0.023	0.010 *	-0.056	0.009 **
Q2	Donut		0.026	0.008 **	0.041	0.026	0.017	0.011	0.029	0.011 *
Q3	ICR	95.4%	-0.043	0.007 **	-0.057	0.017 **	-0.032	0.011 **	-0.044	0.008 **
Q3	Donut		0.008	0.009	0.026	0.027	0.000	0.013	0.003	0.012
Q4	ICR	97.0%	-0.051	0.009 **	-0.039	0.016 *	-0.045	0.015 **	-0.067	0.011 **
Q4	Donut		-0.011	0.011	-0.051	0.027	-0.008	0.016	-0.005	0.017

Notes: Results of quarterly regressions of log consumption change on log change in ICR and donut prices, low-spending individuals only. Individuals with positive consumption in each quarter only. N=919,650 across all years; N=128,412, 388,454, and 402,784 in year pairs 2006-7, 2007-8, and 2008-9, respectively. Proportion of (pooled years) sample for whom initial coverage range (ICR) is marginal in each quarter of first year noted next to estimated coefficients. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

regression that pools all year pairs. High-spending individuals have a large and significant donut price response in quarter 4 (ranging from -0.14 to -0.24 in the individual year pair samples, equalling -0.16 in the pooled analysis) which is significantly larger than the donut response at the beginning of the year (ranging from -0.05 to -0.09 in the individual year pair samples, equalling -0.05 in the pooled analysis). This fact provides striking evidence of myopia, given the low degree of uncertainty that high-spending individuals will be in the donut hole at the end of the year – individuals are on average more than three times as responsive to donut price changes when they are actually *in* the donut hole than they are prior to crossing the donut threshold. Regarding the “spot” price enrollees face prior to the donut hole (the ICR price): in 2006-7 and 2007-8, individuals never significantly respond to the ICR price change, while in 2008-9, the ICR price coefficient is -0.05 to -0.06 at the beginning of the year. The pooled ICR response is significant but small in Q1 and Q2 (-0.02 to -0.03) but shrinks toward zero in Q3-Q4. Each of these patterns is consistent with myopia.

The above results provide estimates of a small, significant price elasticity of demand for prescription drugs throughout the spending distribution. We also observe strong evidence of myopic utilization behavior, in that enrollees’ marginal price response is much more ev-

Table 5: Results of Quarterly ICR and Donut Price Regressions – High-Spending Group

Period	Price	% in Donut	All Years Pooled		2006-7		2007-8		2008-9	
			Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	0.5%	-0.034	0.010 **	-0.018	0.018	-0.012	0.015	-0.057	0.017 **
Q1	Donut		-0.053	0.014 **	-0.054	0.021 **	-0.050	0.024 *	-0.088	0.030 **
Q2	ICR	22.2%	-0.023	0.011 *	-0.003	0.021	-0.013	0.015	-0.053	0.017 **
Q2	Donut		-0.073	0.015 **	-0.077	0.025 **	-0.069	0.022 **	-0.085	0.030 **
Q3	ICR	88.0%	-0.023	0.012	-0.005	0.022	-0.020	0.017	-0.029	0.018
Q3	Donut		-0.086	0.018 **	-0.111	0.029 **	-0.116	0.031 **	-0.069	0.028 *
Q4	ICR	94.9%	-0.008	0.014	-0.021	0.025	-0.017	0.021	0.017	0.022
Q4	Donut		-0.164	0.025 **	-0.167	0.027 **	-0.243	0.047 **	-0.138	0.034 **

Notes: Results of quarterly regressions of log consumption change on log change in ICR and donut prices, high-spending individuals only. Individuals with positive consumption in each quarter only. N=294,898 across all years; N=61,198, 127,052, and 106,648 in year pairs 2006-7, 2007-8, and 2008-9, respectively. Proportion of (pooled years) sample for whom donut hole is marginal in each quarter of first year noted next to estimated coefficients. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

ident at the end of the year, once they have entered their marginal coverage phase. The evidence from high-spending enrollees in particular allows us to reject that individuals are fully forward-looking and to reject that individuals are fully myopic. We can reject forward-lookingness due to the significant increase in price response as more individuals enter the donut hole; we can reject full myopia due to the negative and significant donut hole price response in Q1 (in which 99.5% of high-spending individuals are still in the ICR). Relating these responses to a parameterized model of myopic behavior requires that we explicitly link each quarter’s consumption responses to the distribution of enrollees across coverage phases; we leave that exercise for the next Section. Next, we examine whether individuals respond to “prices” other than the out-of-pocket costs that should be most relevant for them given their observed drug consumption patterns.

5.3 Salience Results

We next estimate the reduced form model with two terms capturing variation in salient coverage characteristics. That is, we use the exact same specification as in Section 5.2, but with two additional variables. The deductible change variable equals -1 if the deductible threshold is decreased between years, and equals 1 if the deductible threshold is increased between years by more than the standard deductible change (e.g., \$250 to \$265 between

2006 and 2007); it equals 0 otherwise.²⁹ The second variable captures changes in categorical coverage in the donut hole. As discussed in Appendix C, we focus on whether consumers respond to nominally large (“stark”) changes in coverage over and above how those changes impact out-of-pocket donut hole price. The variable “Stark” equals 1 if coverage for an entire class of drugs (e.g., all generics) is dropped between year 1 and year 2, and equals -1 if coverage for an entire class of drugs is added; it equals 0 otherwise.

The regression results for the full year are shown in Table 6, which shows that Part D enrollees do respond on average to stark changes in donut hole coverage beyond such changes’ effects on expected donut hole prices. Results are shown separately for high- and low-spending enrollees (for brevity, the regression for all years pooled is shown in the Table). The “Stark” donut hole coverage response is negative and significant for both spending groups, *even among low-spending individuals* who have essentially zero probability of reaching the donut hole. The point estimates indicate that low-spending enrollees, observing that their plan dropped (added) generic or branded coverage to their plan benefit, would decrease (increase) annual consumption by 6% even though they never expect to encounter the donut prices. The “Stark” coverage change response is smaller in magnitude for high-spending individuals, indicating a 2.5% decrease in spending among high-spending enrollees losing gap coverage. It is important to note here that this result is not an artifact of the collinearity between the donut price and the “Stark” variable – the consumption response to the “Stark” variable is similar regardless of whether the specification includes a separate control for donut price (see Appendix Table 14).

We also observe a substantial effect of deductible coverage on consumption. The deductible is not marginal for the vast majority of enrollees, but all enrollees with deductibles spend some portion of the year in that phase, so that the deductible response may be due to myopia (overreacting to the individual-specific effective price change earlier in the year) or to salience (reacting to the deductible based on its visibility in benefit presentation rather

²⁹For plans with any deductible in both years of the year pair, there is essentially no within-phase price instrument variation because the deductible coinsurance equals 100%.

Table 6: Results of Full Year ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – Low- and High-Spending Enrollees, All Years Pooled

Price	Low-Spend Enrollees		High-Spend Enrollees	
	Coef	SE	Coef	SE
ICR	-0.099	0.010 **	-0.026	0.008 **
Donut	0.043	0.014 **	-0.063	0.013 **
Ded. Chg	-0.033	0.009 **	-0.023	0.008 **
Stark	-0.059	0.015 **	-0.025	0.010 *

Notes: Results of full year regressions of log consumption change on log change in ICR and donut prices, as well as changes in stark gap coverage and deductible coverage. All year pairs pooled. N=1,326,301 for low-spending enrollees; N=307,004 for high-spending enrollees. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

than based on implied out-of-pocket price change). The deductible change response is larger in magnitude for low-spending than for high-spending enrollees, but not significantly so.³⁰

Taken together, these results imply that high-level coverage changes (such as changes in deductible or donut generosity) have a large impact on individuals’ behavior. The deductible response may be further evidence of myopia. However, the significant response to stark changes in donut hole coverage goes beyond what we would expect given a rational, forward-looking calculation of what those changes imply for marginal prices – our most striking evidence of price salience is that observed for low-spending individuals, who are very unlikely to encounter the donut hole during the year.

6 Relating Reduced Form to Models of Consumption

Our results thus far have provided evidence of significant price responses overall, substantial myopia, and price salience effects. In order to construct counterfactual estimates for

³⁰Note that we do not interpret these results as indicating a larger response by low-spending enrollees to changing deductible or gap coverage – consumption enters the regression in logs and thus the coefficients for high and low spenders are not directly comparable. For example, if we scale the “Stark” coefficients by year 1 spending, the implied linear response to the “Stark” variable is -44 among low-spending enrollees vs. -51 among high-spending enrollees. In a model with hyperbolic discounting, we would expect the deductible price response to vary inversely with the overall magnitude of spending, but given the standard errors and log specification, we cannot disentangle the effects of myopia and salience using the deductible coefficients. We do observe in Table 7 in the following Section that deductible responses are stronger at the beginning of the year than at the end of the year, which is stronger evidence that the deductible response is driven in part by myopia.

individuals outside our regression sample (i.e., near the donut hole kink), we require a more general model of consumption behavior. Traditional candidates for this model would include marginal price response and average price response. Dynamic models of consumption along the lines of Aron-Dine et. al. 2015 or Dalton, Gowrisankaran, and Town (2015) allow beneficiaries to respond either to the current spot price or the end of year marginal price. Motivated by our findings in Section 4.2, we also allow consumers to respond to current or future prices, and we allow price salience to factor into consumption patterns. Our reduced form parameters tell us how enrollees in-sample respond to price changes at different points in the budget set; the exercise in this section relates the reduced form price coefficients to structural parameters that capture price sensitivity, myopia, and salience. The relationship between the structural parameters and the reduced form coefficients is mediated by the distribution of enrollees across the different budget set segments at different points in time. After estimating our dynamic structural model, we compare our results to marginal and average price models as well as models which allow for a myopic response to current prices or marginal prices but do not allow for price salience.

The dynamic “current-future” model with salience fits the data well both inside and outside the estimation sample and outperforms alternative models – both the rational marginal price response model and the behavioral average price response model. The model performs well even though it is estimated using linear consumption responses; in Appendix E.2, we derive and estimate a structural model with bunching and coverage phase switching and show that the performance of the richer model is weakly worse than the simpler specification. We then use the preferred model to simulate the effects of filling in the donut hole.

6.1 Parameter Estimates

Inferring a set of structural parameters from our reduced form evidence requires comparing the estimated coefficients across individuals, coverage phases, and points in time. The estimates from our log-log regression specifications are not directly comparable, as both ICR and

donut prices and high and low-spending individuals' consumption levels differ in scale – thus, this exercise requires a simple transformation. Specifically, we first differentiate the log-log model to determine what it implies for (locally) linear demand (recall that the model in Section 4.1 assumed linear demand). Our main (dynamic) empirical specification, including the deductible and “Stark” donut hole coverage terms, is the following:

$$\begin{aligned}
\log(q_{ity}) - \log(q_{it,y-1}) &= a_t + X_{it} * \delta_t + b_{ICR,t} * (\log(P_{ICR,y}) - \log(P_{ICR,y-1})) \\
&+ b_{Donut,t} * (\log(P_{Donut,y}) - \log(P_{Donut,y-1})) \\
&+ \theta_{Ded,t} * dedchg_i + \theta_{stark,t} * stark_i + u_{ity}
\end{aligned}$$

The structural model we wish to link the empirical specification to is:³¹

$$q = \alpha + \eta(\beta MP + (1 - \beta)CP) + K_d * \mathbb{1}(dedchg = 1) + K_s * \mathbb{1}(stark = 1).$$

In this model, the term α is predicted spending (in year 2) at zero prices and absent any deductible or donut hole coverage change between year 1 and year 2. The static price response is η – this is the price response we would observe under a linear price contract. The hyperbolic discount factor is β . K_d and K_s capture consumption shifts between years 1 and 2 in response to changes in the deductible and stark donut hole coverage, respectively, which we model as distinct from the budget set responses captured by η and β .

Letting \mathbf{z}_y be the vector of year y prices, we linearize the specification around \mathbf{z}_1 . Details are given in Appendix D. This gives us an expression for each of the below coefficients from our structural model (omitting subscripts for the sake of brevity). These relations show that the linear structural constant α varies with year 1 consumption and enrollee characteristics, so that this exercise yields a *local* linearization of enrollees' expected year 2 consumption as a function of year 1 observed consumption and predicted year-to-year trends given observables.

³¹In the baseline estimation, we ignore uncertainty, bunching, and coverage phase switching, each of which is likely to have limited impact in our estimates given that regression sample enrollees were chosen to be well away from coverage thresholds. The robustness exercise in Appendix E.2 relaxes these assumptions and the estimates show that results are unchanged.

The ICR and donut price coefficients vary as a function of observables as well as the ratio of year 1 consumption to each respective year 1 price.

$$\alpha = q_1 * \exp(a + X * \delta + u)(1 - b_{ICR} - b_{Don}) \quad (4)$$

$$\eta * (\beta * \Pr(MP = ICR) + (1 - \beta) * \Pr(CP = ICR)) = \left(\frac{q_1}{P_{ICR,1}} \right) * \exp(a + X * \delta + u) * b_{ICR} \quad (5)$$

$$\eta * (\beta * \Pr(MP = Don) + (1 - \beta) * \Pr(CP = Don)) = \left(\frac{q_1}{P_{Don,1}} \right) * \exp(a + X * \delta + u) * b_{Don} \quad (6)$$

$$K_d = q_1 * \exp(a + X * \delta + u) * \theta_{Ded} \quad (7)$$

$$K_s = q_1 * \exp(a + X * \delta + u) * \theta_{stark} \quad (8)$$

Table 7 below displays the results of our richest reduced form specification, allowing for both myopia and salience (using the same identification strategy as in all results in Section 5.2): we regress year to year quarterly consumption changes on ICR and donut price changes as well as the deductible change and “Stark” variables, for high and low-spending enrollees only. We allow ICR and donut price responses to vary by spending group, as the proportion of individuals for whom each phase is marginal or current in each period will differ across groups; we hold all other coefficients fixed across groups. These are the full set of reduced form results we use to infer structural parameters.

The results shown in Table 7 are consistent with the patterns described in Section 5.2: low-spending enrollees’ marginal price (ICR price) response is flat over the year and high-spending enrollees’ marginal price (donut price) response is steeply increasing in magnitude over the year. Enrollees’ deductible response is decreasing in magnitude over the course of the year (this is consistent with myopia as enrollees only encounter deductible prices early in the year). Finally, the response to the “Stark” variable is non-monotonic over the year – it is largest in Q1 and Q4.³²

³²This feature of the results is driven by low-spending enrollees’ non-monotonic response to stark changes in donut hole coverage – as we see in Appendix Table 16, which compares the results in Table 7 to the same results obtained from separate regressions for the high- and low-spending samples, low-spending enrollees’ response to the “Stark” variable is negative and significant in Q1 and Q4 (the Q4 response is larger, but not significantly so), while high-spending enrollees’ response to the “Stark” variable is only significant in

Table 7: Results of Quarterly ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – Low- and High-Spending Enrollees, All Years Pooled

Period	Price	All Years Pooled (Pooled Regression)			
		Low-Spending Enrollee Response		High-Spending Enrollee Response	
		Coef	SE	Coef	SE
Q1	ICR	-0.053	0.006 **	-0.026	0.008 **
Q1	Donut	0.019	0.007 **	-0.048	0.010 **
Q1	Ded. Chg	-0.049	0.005 **	-0.049	0.005 **
Q1	Stark	-0.014	0.009	-0.014	0.009
Q2	ICR	-0.045	0.006 **	-0.011	0.009
Q2	Donut	0.026	0.007 **	-0.044	0.013 **
Q2	Ded. Chg	-0.024	0.005 **	-0.024	0.005 **
Q2	Stark	0.002	0.009	0.002	0.009
Q3	ICR	-0.043	0.006 **	-0.003	0.009
Q3	Donut	0.012	0.008	-0.083	0.012 **
Q3	Ded. Chg	-0.016	0.006 **	-0.016	0.006 **
Q3	Stark	-0.007	0.009	-0.007	0.009
Q4	ICR	-0.056	0.008 **	0.024	0.012 *
Q4	Donut	0.003	0.010	-0.168	0.016 **
Q4	Ded. Chg	-0.012	0.007	-0.012	0.007
Q4	Stark	-0.047	0.011 **	-0.047	0.011 **

Notes: Results of quarterly regressions of log consumption change on log change in ICR and donut prices, as well as changes in stark gap coverage and deductible coverage. Regression for high- and low-spending enrollees, with separate ICR and donut responses for each spending group. N=1,214,548. "Stark" gap coverage and deductible change response held fixed across spending groups. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

We use the results of the above regression to infer our structural model parameters. We pool price response estimates across groups of observations (individual-quarters) whose price responses are expected to be similar by classifying sample observations by quintile of $\left(\frac{q_{it,1}}{P_{i,1}}\right) * \exp(\hat{a}_t + X_i * \hat{\delta}_t + \hat{u}_{it,2})$, for each price P_i , giving us 25 groups overall – this classification allows us to define groups of individuals based on similarity in their expected price changes, both in and outside the regression sample. We then estimate a single η_g for each group g .³³ We impose a single discount factor β across all sample individuals. Using the expressions including the parameters η and β in the above Taylor expansions (equations (2) and (3) above), we use a GMM procedure to estimate our 25 static linear price response parameters η and our hyperbolic discount parameter β . Equations (2) and (3) require $\Pr\{CP = ICR\}$, $\Pr\{MP = ICR\}$, $\Pr\{CP = Donut\}$, and $\Pr\{MP = Donut\}$ as inputs – for the estimates

Q4. These results are consistent with the low-spending individuals' response to stark donut hole coverage being driven by salience effects and the high-spending individuals' response being driven at least in part by myopia.

³³The error term $u_{it,2}$ is not directly observed, so we use Duan's smearing technique to scale all transformed coefficients based on the distribution of the regression residuals. We allow for heteroscedasticity and let the smearing factor vary in demographic variables.

reported, we assumed “perfect foresight” regarding current and marginal coverage phases, in that we used the actual observed probabilities of each phase being current or marginal in year 2 for each individual and each quarter.³⁴

It is worth noting here that, even after allowing marginal price responses to vary across 25 groups of individuals, we are overidentified. We are inferring 28 structural parameters from 32 reduced form price responses, pooling information across individuals, coverage phases, and points in time. No simple structural model is expected to fit all coefficients exactly; for example, none of the models under consideration would predict the small positive donut price responses among low-spending individuals. The goal instead is to determine which model of behavior best fits all the observed data and to compare performance across models.

Counterfactual simulation requires that we extrapolate the structural parameters for individuals outside the regression sample (high and low spending enrollees). We extrapolate α , K_s , and K_d using equations (1), (4), and (5) above. We obtain η and β from the GMM procedure – individuals outside the regression sample falling in group g are assigned the structural static price response $\hat{\eta}_g$ and discount factor $\hat{\beta}$. In order to ensure that we extrapolate only to individuals who are similar to our regression sample, we exclude outliers, defined as individuals whose $\exp(X_{it} * \hat{\delta}_t + \hat{u}_{it,2})$ lie below the 1st percentile or above the 99th percentile of the same metric in the regression sample.

Table 8 displays the mean values for all model parameter estimates for individuals in and out of the regression sample. At zero prices, our sample is predicted to consume 371 days supply of drugs per quarter. The average linear price response $\eta = -1.7$ implies an average static price elasticity of -0.13 (evaluated at marginal prices). The quarterly hyperbolic discount factor $\beta = 0.31$ suggests substantial myopia – it implies that, prior to entering

³⁴Of course, enrollees’ expectations could instead be that they will consume in year 2 exactly as they did in year 1, or they could expect that they will be a random draw from the observed distribution of consumption among *similar* individuals based on year 1 consumption. Our estimates are not sensitive to this assumption – re-estimating the structural parameters using the individual’s actual year 1 phase probabilities or “rational expectations” phase probabilities generates similar estimates, as discussed with regard to model fit below. In the rational expectations case, we classify individuals in each year pair into 100 cells by centiles of year 1 total spending and run the year 2 claims of 200 persons in each cell through the cost parameters for the plan for each individual in the cell and take the means of the resulting phase probabilities.

Table 8: Estimated Structural Model Parameters, All Years Pooled

Description	Structural Parameter	Central Estimate	Standard Error
Days supply (P=0)	α	370.864	13.746
Myopia	β	0.312	0.081
Marginal price effect	η	-1.656	0.146
Deductible effect	K_{deduct}	-8.833	1.666
Stark gap effect	K_{stark}	-6.335	2.698
Implied elasticity	ϵ	-0.127	0.011

Notes: Authors' calculations. Data are shown only for individuals within the 1st to 99th percentiles of the distribution of the predicted trend in consumption; individuals too dissimilar from the regression sample in this dimension are dropped. All parameters shown are averages except for the hyperbolic discount factor. Standard errors from nonparametric bootstrap of entire procedure, 200 iterations.

their marginal coverage phase, individuals are more than twice as responsive to the price they currently face as they are to their marginal price. Finally, on average, eliminating the deductible or adding stark donut hole coverage is predicted to increase consumption per quarter slightly, by 9 and 6 days supply, respectively.

The static price elasticity of -0.13 estimated using this procedure is similar in magnitude to the literature (see, e.g., Chandra, Gruber, and McKnight (2010)), but is notably smaller than the static elasticities estimated in Einav, Finkelstein, and Schrimpf (2015) in the Part D context – they estimate elasticities ranging from -0.3 to -0.5 . Their identification strategy infers price sensitivity from the magnitude of bunching behavior observed at the Part D kink and is thus focused on a small sample we explicitly exclude from our regression analysis – for this reason, we may not expect to arrive at the same behavioral elasticity. However, as we note below, the elasticity obtained from our reduced form procedure predicts bunching at the kink well in a dynamic model.

Our estimate of myopia, based on variation from enrollees across the spending and age distributions and price changes at different points in the budget set, is between those estimated in two recent papers using narrower variation. Einav, Finkelstein, and Schrimpf (2015) estimate a weekly exponential discount factor of 0.96 using variation in timing of enrollees joining Part D plans; if we assumed that this model of discounting were true, it would imply that an individual discounts all future quarters by an average factor of 0.49

across Q1-Q3, implying significantly more forward-looking consumption behavior than we estimate. In contrast, Dalton, Gowrisankaran, and Town (2015), relying on identification from behavior changes right at the donut kink, find evidence of full myopia.

In order to investigate the behavioral determinants of myopia and salience in the results in Table 8, we repeated our entire estimation procedure separately for individuals with and without chronic conditions, and for groups of individuals defined by specific chronic conditions.³⁵ This allows us to examine whether individuals taking maintenance drugs on a predictable basis (as opposed to taking drugs for acute conditions) exhibit the same deviations from rationality as our overall sample. The results comparing the chronically ill to the non-chronically ill and the results for the most popular chronic conditions in our sample (hypercholesterolemia, hypertension, and diabetes) are shown in Appendix Table 17. The estimated hyperbolic discount factor β is larger in magnitude for the chronically ill (particularly among hypercholesterolemics and diabetics, who exhibit β s of 0.46 and 0.42, respectively) than for the non-chronically ill, which is consistent with the chronically ill being more forward-looking. However, the most striking feature of the table is the consistency in the parameter estimates across groups – even those with chronic conditions are substantially more sensitive to marginal prices after they encounter those prices than earlier in the year.

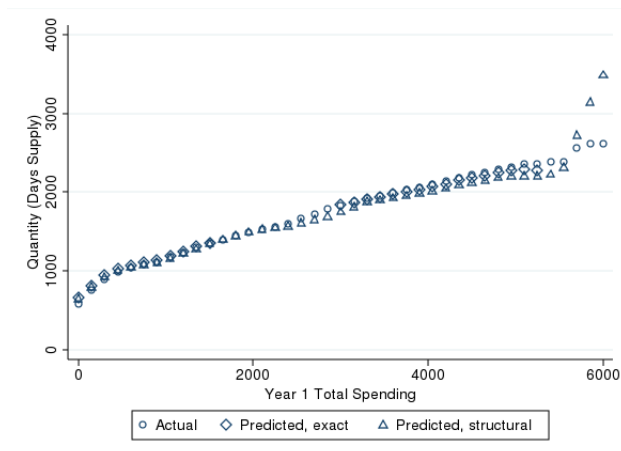
Using the model from Section 4 and the estimated structural parameters from Table 8, we solve for the optimal dynamic consumption path in response to the full nonlinear budget set for all sample individuals, including those near budget set kinks. Figure 2 displays actual and predicted consumption (summed over the full year) for individuals throughout the year 1 spending distribution.³⁶ Two versions of the prediction are displayed – the exact quantity predicted by the regression model, in sample, and the simulated results using the structural model parameters. Both predictions work quite well for individuals in the regression sample (high and low-spending individuals) – the structural model under predicts actual year 2

³⁵Following Goldman, et al. (2004), we identify seven chronic illnesses – hypercholesterolemia, hypertension, diabetes, gastritis, arthritis, asthma, and affective disorders – using diagnosis codes from the individuals’ medical claims histories.

³⁶Predicted consumption is simulated consumption throughout the year given actual year 2 prices.

spending by 0.34% on average. Notably, the structural model also replicates consumption for enrollees near the donut hole threshold quite well – we over predict year 2 spending by 0.18% on average for this sample. The prediction performs less well as we approach the very

Figure 2: Actual and Predicted Year 2 Quantity Consumed – Outliers Excluded



high part of the spending distribution and is poor for the small number (1% of the sample) of non-outliers above the top of the “high” spending group (the \$5,000 cutoff).³⁷ In the counterfactual simulations, individuals above the “high-spending” cutoff are excluded.³⁸

The absolute and relative distribution of actual and predicted spending are shown in Figure 3.³⁹ Predicted spending is close to actual spending on average, as implied by the comparison in Figure 2. The left panel of the Figure shows that we slightly underpredict low spending and overpredict high spending. The right panel shows the distribution of actual and predicted spending relative to the donut threshold (recall that donut thresholds increase in each year of the sample). Even though we estimated price sensitivity using only

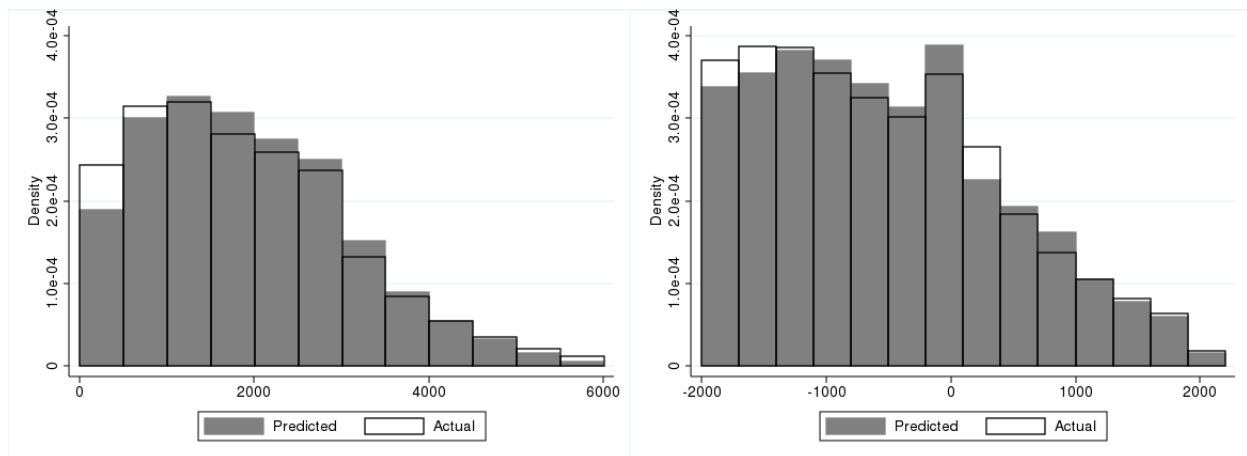
³⁷The small 1% sample of non-outliers in the above-\$5,000 range accounts for 3% of non-outliers’ spending. In the full sample, those spending above \$5,000 account for 6% of enrollees and 18% of overall spending, as there is a very long tail to the expenditure distribution. Lack of overlap in covariates between the full range of very high-spending enrollees and our regression sample limits our ability to extrapolate price responses to that part of the distribution.

³⁸Results are similar whether the perfect foresight approach is used as the model of expectations in the GMM estimation or whether, alternatively, we use the individual’s year 1 consumption (0.3% error) or rational expectations (-0.4% error).

³⁹Here, because coverage thresholds are defined based on total drug spending rather than quantity consumed, we compare actual and predicted spending rather than consumption as in Figure 2.

individuals away from the donut threshold, we are able to replicate bunching at the donut kink quite well. The slight over prediction of bunching behavior exactly at the kink is due to our assumption of no uncertainty in the simulation model – allowing for uncertainty would yield some dispersion in excess mass around the kink, as we observe in the actual spending distribution just to the right of 0.

Figure 3: Distribution of Actual and Predicted Year 2 Spending – Outliers Excluded
(a) Full Distribution – All Years **(b)** Distribution Relative to Donut Threshold



The above reduced form and structural analyses showed evidence of imperfectly forward-looking behavior, which rejects a model of consumers responding rationally to marginal price. However, the structural model we estimated above, while intuitive, could miss other behavioral consumption patterns, such as response to an average price as has been found in the empirical literature on electricity consumption. A comparison of the full structural model (with myopia and salience terms) to two alternative models – an average price response model and a marginal price response model – is shown in Table 9. The Table shows that the structural model (0.18% error) performs better in predicting out-of-sample spending than either the average price (6.82% error) or marginal price model (10.34% error). Similarly, the mean-squared error in prediction (MSE) is much smaller for the structural model than either alternative, both in- and outside the regression sample.

Table 9: Actual and Predicted Year 2 Consumption and Spending – Structural Model, Average Price Model, and Marginal Price Model

	Structural Model Comparison				Average Price Model Comparison				Marginal Price Model Comparison			
	Actual	Predicted		MSE	Actual	Predicted		MSE	Actual	Predicted		MSE
		Mean	% diff			Mean	% diff			Mean	% diff	
In-Sample (Low/High Spending Enrollees)	1,718.239	1,703.561	-0.85%	996,251	1,755.345	1,799.549	2.52%	1,726,621	1,716.097	1,810.151	5.48%	1,705,993
Out-of-Sample (Medium Spending Enrollees)	2,324.124	2,328.278	0.18%	990,186	2,320.353	2,478.690	6.82%	1,816,324	2,316.501	2,555.971	10.34%	1,915,258
All Spending Groups	1,972.520	1,965.746	-0.34%	993,706	1,993.213	2,085.467	4.63%	1,764,386	1,963.783	2,117.827	7.84%	1,792,322

Notes: Comparison of actual and simulated spending, structural model vs. average price model and marginal price model. Average and marginal price models estimated using quarterly regressions of log consumption change on year-to-year change in average price and marginal price, respectively. The same controls as in Section 6 are used in the regressions and the same smearing technique as in Section 7 is employed to transform predictions. In each comparison, data are shown only for individuals within the 1st to 99th percentiles of the distribution of the predicted trend in consumption; individuals too dissimilar from the regression sample in this dimension are dropped. Average and marginal prices determined by applying perfect foresight weights to coverage phase-specific individual prices.

6.2 Counterfactual – Filling in the Donut Hole

We next use the estimated structural parameters to simulate the effect of filling in the donut hole on total spending for low and high-spending individuals as well as for individuals near the donut hole kink. We impose that “filling in the donut hole” takes the form of setting the donut hole price in each plan equal to the ICR price (so that there is no donut kink). For all individuals with no donut hole coverage or generic only donut hole coverage in year 1 of the relevant year pair, we also set the “Stark” variable equal to -1 , as filling in the donut hole would entail a stark increase in donut hole coverage.⁴⁰ Table 10 shows the effect of filling

Table 10: Estimated Effect of Filling in the Donut Hole, Pooled All Years

Estimated impact of filling in the donut hole	\$114.17	5.81%
Based on price response alone	\$77.25	3.93%
Additional effect of price salience	\$36.92	1.88%

Notes: Authors' calculations. Data are shown only for individuals within the 1st to 99th percentiles of the distribution of the predicted trend in consumption; individuals too dissimilar from the regression sample in this dimension are dropped.

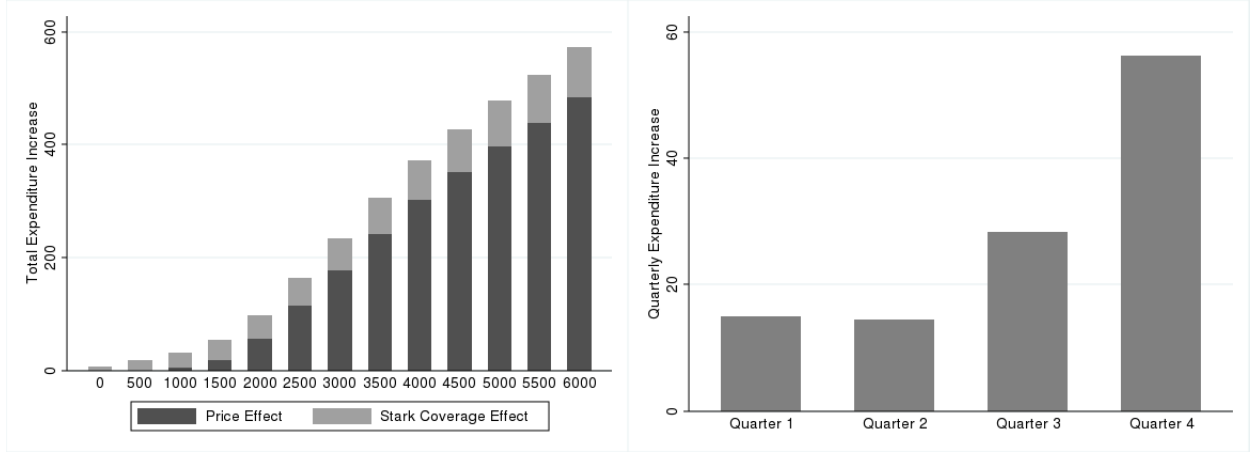
⁴⁰An important caveat is necessary here. This counterfactual is performed (1) relative to observed year 2 prices in each sample year pair; and (2) holding all other plan features fixed. Thus, we should interpret the results as capturing *what 2007, 2008, and 2009 consumption would have been if each plan's donut hole were filled in, holding enrollment and other features of plan generosity fixed*. In order to determine the effect of changes to the donut hole generosity between now and 2020, one would need to account for trends in generosity between our sample period and the present, as well as how plans adjust other plan features in response to the imposed change in donut generosity (which may impact, in turn, sorting of patients across plans). Without a model to inform supply side responses to the ACA's donut policy changes, we leave the more sophisticated counterfactual to future research.

in the donut hole on mean spending, in levels and percentages. Filling in the donut hole would increase spending by \$114 on average, or 6%. \$77 of this increase is due to the price response (setting the donut price equal to the ICR price), but \$37 (or 32% of the overall effect) comes from what we call the “salience” effect, the coefficient on the stark increase in donut hole coverage.

Figure 11 shows heterogeneity in the effect of filling in the donut. The left panel shows how the effect of filling in the donut varies with position in the spending distribution. That is, we plot the mean increase in spending when we fill in the donut vs. year 1 spending. The price and salience effects are shown separately. The price effect is monotonically increasing in the magnitude of spending – individuals whose marginal price is the donut hole price are impacted more by the price change than individuals who do not hit the donut. Moreover, individuals who consume more of their prescription drugs while in the donut hole (higher spenders within the donut-marginal group) are more affected by the price change due to myopia. On the other hand, the “Stark” or what we called the “salience” effect is present throughout the spending distribution, so that even low spending individuals are expected to increase spending in response to the donut hole being filled in. Notably, much of the effect of filling in the donut hole is due to price responses for individuals who end the year well away from the kink and who would accordingly not exhibit bunching behavior in analyses that focus on the kink. For example, observations for enrollees ending a given year at a spending level within \$200 of the donut threshold appropriate for that year account for only \$16 (14%) of the \$114 average effect of filling in the donut hole.

The right panel of Figure 11 plots the mean increase separately by quarter. As implied by our estimate of the hyperbolic discount factor $\beta = 0.32$, individuals are much less responsive to filling in the donut hole at the beginning of the year than they are at the end of the year.

Figure 11: Heterogeneity in Effect of Filling in the Donut Hole
(a) (b)



7 Discussion

We examine prescription drug consumption in the context of Medicare Part D, an insurance program in which enrollees are exposed to substantial cost-sharing incentives and in which nonlinear, complex price schedules are found to lead to additional price responses beyond those anticipated by the designers.

Our identification strategy allows us to estimate static and dynamic price elasticities using consumption responses to price variation in multiple regions of the nonlinear budget set. Part D enrollees are very unlikely to switch plans between years, giving us a plausibly exogenous source of identifying variation in the form of year-to-year changes in plan generosity for individuals already enrolled in Part D plans. Demand models under nonlinear budget sets can be challenging to estimate; in the absence of enough variation to estimate fully nonparametric responses, the solution is typically to specify a structural model of consumer optimization. In order to accommodate multiple behavioral models of consumption suggested by the nonlinear budget set literature in other contexts, we begin by using linear methods to estimate reduced form consumption responses to prices throughout the budget set. We focus on individuals whose marginal prices are highly likely to be in the interior of budget set segments in order to minimize the impact of nonlinear responses; analysis in the Appendix demonstrates that our results are robust to further sample restrictions and more complex modeling. We use

the reduced form patterns to estimate a structural model with imperfect forward-looking behavior and price salience effects, and demonstrate that the model fits the data well; it outperforms both a fully rational model in which consumers respond to marginal price, as well as an alternative behavioral model in which consumers respond to average price. Notably, the model performs quite well outside the regression sample, which allows us to simulate consumption responses to counterfactual price schedules for enrollees throughout the spending distribution, including those near the donut hole kink.

We find that, while enrollees' static price elasticities are of a similar magnitude to estimates in the prior literature, Part D enrollees also exhibit certain behavioral biases in their consumption patterns related to the structure of Part D cost-sharing. In particular, we demonstrate evidence of imperfectly forward-looking behavior, in that enrollees are much more responsive to cost-sharing in current periods than in future periods. We also find that enrollees respond to more salient plan benefit changes, such as addition or removal of entire categories of drugs from donut hole coverage, beyond how those changes impact enrollees' actual out-of-pocket prices. Given that rational optimization of consumption in a setting such as Part D requires a complicated calculation with many inputs, these results may not be surprising.

The results described above yield some striking insights into insurance enrollees' decision-making, and carry important implications for policy. We find that, all else equal, filling in the Part D donut hole will increase spending by \$114, or 6%, for the average enrollee in our sample. Over 30% of this effect is due to salience and, accordingly, impacts even relatively low-spending Part D enrollees. The remainder of the effect occurs primarily at the end of the year, due to imperfect forward-looking behavior, and falls disproportionately on higher-spending enrollees who are more likely to enter the donut hole.

Our findings suggest that prescription drug plan designers must carefully account for consumers' dynamic incentives and understanding of complex price schedules. The welfare implications of the current Part D plan design will be a function of the price responses

documented here as well as the health impacts of altering drug consumption in response to prices and the overall program costs in- and outside Part D. We leave these topics for future research. This paper also developed a useful methodology for analyzing consumption in complex, nonlinear environments by estimating structural parameters using variation in linear regions of the budget set; we hope to extend the methodology to other applications, such as income taxation, in future work.

8 References

- Abaluck, J. and J. Gruber. 2011. "Heterogeneity in Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review*. May. 101(3) 377-381.
- Agarwal, S., J. Driscoll, X. Gabaix, and D. Laibson. 2009. "The Age of Reason: Financial Decisions over the Life-Cycle with Implications for Regulation." *Brookings Papers on Economic Activity*. October.
- Aron-Dine, A., L. Einav, A. Finkelstein, and M. Cullen. 2015. "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?" *Review of Economics and Statistics*, forthcoming.
- Blomquist, W. and W. Newey. 2002. "Nonparametric Estimation with Nonlinear Budget Sets." *Econometrica: Journal of the Econometric Society*. November. 70(6) 2455-2480.
- Brot-Goldberg, Z., A. Chandra, B. Handel, and J. Kolstad. 2015. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." Mimeo. Available at http://faculty.haas.berkeley.edu/jkolstad/Utilization_BCHK_Web.pdf.
- Centers for Medicare and Medicaid Services (CMS). 2015. "Closing the Coverage Gap – Medicare Prescription Drugs Are Becoming More Affordable." *CMS Product No. 11493*. January. Available at <https://www.medicare.gov/Pubs/pdf/11493.pdf>.
- Chandra, A., J. Gruber, and R. McKnight. 2010. "Patient Cost-Sharing and Hospitalization Offsets in the Elderly." *American Economic Review*. March. 100(1) 193-213.
- Chetty, R., A. Looney, and K. Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review*. 99(4): 1145-77.
- Dalton, C. M., G. Gowrisankaran, and R. Town. 2015. "Myopia and Complex Dynamic Incentives: Evidence from Medicare Part D." Mimeo. Available at http://users.wfu.edu/daltonc/docs/Research/Dalton_Gowrisankaran_Town_PartD.pdf.
- Duan, N. 1983. "Smearing Estimate: A Nonparametric Retransformation Method." *Journal of the American Statistical Association* 78(383) 605-610.
- Duggan, M., P. Healy, and F. Morton. 2008. "Providing Prescription Drug Coverage to the Elderly: America's Experiment with Medicare Part D." *Journal of Economic Perspectives*. Fall. 22(4) 69-92(24).
- Einav, L., A. Finkelstein, and P. Schrimpf. 2015. "The Response of Drug Expenditures to Non-Linear Contract Design: Evidence from Medicare Part D." *Quarterly Journal of Economics*. May. 130(2) 841-899.
- Ericson, K. 2014. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy*. February. 6(1): 38-64.
- Feldman, N., P. Katuscak, and L. Kawano. 2015. "Taxpayer Confusion: Evidence from the Child Tax Credit." Mimeo. July. Available at SSRN: <http://ssrn.com/abstract=2296542> [accessed August 27, 2015].
- Fratiglioni, L., D. De Ronchi, and H. Aguero-Torres. 1999. "Worldwide Prevalence and Incidence of Dementia." *Drugs and Aging*. November. 15(5) 365-375.
- Goldman, D., G. Joyce, and Y. Zheng. 2007. "Prescription Drug Cost Sharing." *Journal of the American Medical Association*. July. 298(1) 61-69.

- Goldman, D. P., G. F. Joyce, J. J. Escarce, J. E. Pace, M. D. Colomom, M. Laouri, P. B. Landsman, and S. M. Teutsch. 2004. "Pharmacy Benefits and the Use of Drugs by the Chronically Ill." *Journal of the American Medical Association* 291, 2344-23501.
- Handel, B. and J. Kolstad. 2015. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review*. August. 105(8), 2449-2500.
- Hausman, J. 1985. "The Econometrics of Nonlinear Budget Sets." *Econometrica: Journal of the Econometric Society*. November. 53(6) 1255-1282.
- Heiss, F., D. McFadden, and J. Winter. 2006. "Who Failed to Enroll in Medicare Part D, and Why?" *Health Affairs*. September-October. 25(5) w344-354.
- Ho, K., J. Hogan, and F. Scott Morton. 2015. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." *National Bureau of Economic Research Working Paper No. 21028*. March.
- Ito, K. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*. February. 104(2) 537-563.
- Ketcham, J., C. Lucarelli, E. Miravete, and M. Roebuck. 2012. "Sinking, Swimming, or Learning to Swim in Medicare Part D." *American Economic Review*. October. 102(6) 2639-2673.
- Ketcham, J. and K. Simon. 2008. "Medicare Part D's Effects on Elderly Patients' Drug Costs and Utilization." *American Journal of Managed Care*. 14(11 Suppl): 14-22.
- Kling, J., S. Mullainathan, E. Shafir, L. Vermeulen, and M. Wrobel. 2008. "Misperception in Choosing Medicare Drug Plans." *Unpublished manuscript*. November.
- Kowalski, A. 2015. "Estimating the Tradeoff Between Risk Protection and Moral Hazard with a Nonlinear Budget Set Model of Health Insurance." *National Bureau of Economic Research Working Paper No. 18108*. July.
- Laibson, D. A. Repetto, and J. Tobacman. 2015. "Estimating Discount Functions with Consumption Choices over the Lifecycle." *Mimeo*. July. Available at <http://assets.wharton.upenn.edu/tobacman/papers/Estimating%20Discount%20Functions.pdf>.
- Lichtenberg, F. and S. Sun. 2007. "The Impact of Medicare Part D on Prescription Drug Use by the Elderly." *Health Affairs*. November-December. 26(6) 1735-1744.
- Liebman, J. B. and R. J. Zeckhauser. 2004. "Schmeduling." October.
- Loewenstein, G., et al. 2013. "Consumers' Misunderstanding of Health Insurance." *Journal of Health Economics* 32: 850-862.
- Manning, W., J. Newhouse, N. Duan, E. Keeler, and A. Leibowitz. 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review*. June. 77(3) 251-277.
- Salthouse, T. 2004. "What and When of Cognitive Aging." *Current Directions in Psychological Science*. 13(4) 140-144.
- Yin, W., A. Basu, J. Zhang, A. Rabbani, D. Meltzer, and C. Alexander. 2008. "The Effect of the Medicare Part D Prescription Benefit on Drug Utilization and Expenditures." *Annals of Internal Medicine* 148 169-177.

Appendix

A Price and Instrument Construction

In order to illustrate how our prices and price instruments are calculated, consider the following example. Suppose that, in 2006, the individual in question takes two drugs monthly, drug X and drug Y; in 2007, the individual also takes drug Z. In 2006, the individual is enrolled in plan A; in 2007, she switches to plan B. The retail prices and cost-sharing for Plans A and B, drugs X, Y, and Z, and years 2006 and 2007 are shown in Appendix Table 8. As we see in the Table, plan A has coverage of generics (drug Y only) in the donut hole in both years, while drug B has no donut hole coverage in either year. In general, both retail prices and copays are different across plans for each drug and across years for each plan-drug.

Appendix Table 8: Retail Prices and Out-of-Pocket Costs for Example Plans and Drugs

Drug	2006 Retail and Out-of-Pocket Prices					2007 Retail and Out-of-Pocket Prices				
	Plan A, 2006					Plan A, 2007				
	<i>Retail Price</i>	<i>Deductible</i>	<i>ICR</i>	<i>Donut Hole</i>	<i>Catastrophic</i>	<i>Retail Price</i>	<i>Deductible</i>	<i>ICR</i>	<i>Donut Hole</i>	<i>Catastrophic</i>
X	100.00	100.00	30.00	100.00	5.00	110.00	110.00	31.00	110.00	5.50
Y	30.00	30.00	10.00	10.00	2.00	32.00	32.00	12.00	12.00	2.15
Z	150.00	150.00	45.00	150.00	7.50	160.00	160.00	45.00	160.00	8.00
	Plan B, 2006					Plan B, 2007				
	<i>Retail Price</i>	<i>Deductible</i>	<i>ICR</i>	<i>Donut Hole</i>	<i>Catastrophic</i>	<i>Retail Price</i>	<i>Deductible</i>	<i>ICR</i>	<i>Donut Hole</i>	<i>Catastrophic</i>
X	115.00	115.00	40.00	115.00	5.75	117.00	117.00	40.00	117.00	5.85
Y	25.00	25.00	12.00	25.00	2.00	27.00	27.00	8.00	27.00	2.15
Z	130.00	130.00	50.00	130.00	6.50	130.00	130.00	55.00	130.00	6.50

The corresponding phase-year-specific prices and instruments are shown for each drug and on average across all drugs the individual takes in Appendix Table 9. Recall that our procedure requires a single price and instrument for each individual-year. For the “actual” prices, we aggregate the prices in plan A in 2006 using 2006 weights to obtain the 2006 average prices (e.g., the ICR price in 2006 is \$20, the average copay in plan A across drugs X and Y) and we aggregate the prices in plan B in 2007 using 2007 weights to obtain the 2007 average prices (e.g., the donut price in 2007 is \$91.33, the average copay in plan B across drugs X, Y, and Z). However, to obtain the instruments, we hold plan choice, retail price, and consumption weights fixed at 2006 values. Hence, the price instrument in 2006 is the same as the actual price in 2006 in each phase, and the 2007 price instruments differ from the 2006 price instruments only insofar as plan A’s generosity changed between years 2006 and 2007 holding retail price fixed. In this example, there is no deductible price change in the instrument for any drug or on average, the 2006-2007 price change in the ICR equals the change in plan A’s copays 2006-2007, and the donut price only changes for drug Y, the drug with some donut coverage in plan A.

Appendix Table 9: Average Prices and Price Instruments for Example Individual

	Drug	2006 Phase Prices and Instruments					2007 Phase Prices and Instruments				
		Quantity	Deductible	ICR	Donut Hole	Catastrophic	Quantity	Deductible	ICR	Donut Hole	Catastrophic
"Actual" Prices	X	30.00	100.00	30.00	100.00	5.00	30.00	117.00	40.00	117.00	5.85
	Y	30.00	30.00	10.00	10.00	2.00	30.00	27.00	8.00	27.00	2.15
	Z	0.00	150.00	45.00	150.00	7.50	30.00	130.00	55.00	130.00	6.50
	Average		65.00	20.00	55.00	3.50		91.33	34.33	91.33	4.83
Price Instruments	X	30.00	100.00	30.00	100.00	5.00	30.00	100.00	31.00	100.00	5.35
	Y	30.00	30.00	10.00	10.00	2.00	30.00	30.00	12.00	12.00	2.15
	Z	0.00	150.00	45.00	150.00	7.50	0.00	150.00	45.00	150.00	7.50
	Average		65.00	20.00	55.00	3.50		65.00	21.50	56.00	3.75

B Comparison of OLS, Reduced Form, and IV Regressions

Appendix Table 10 shows a comparison of the ordinary least squares, reduced form, and instrumental variables regressions of log quarterly change in consumption on log price (or instrumental variable price) change in the ICR and donut regions of the budget set. The IV results are broken out by year in Tables 4 and 5 in the main text.

Appendix Table 10: Results of Quarterly ICR and Donut Price Regressions – Low-Spending and High-Spending Groups

Period	Price (IV)	Ordinary Least Squares		Reduced Form		Instrumental Variables	
		Coef	SE	Coef	SE	Coef	SE
Q1	ICR	0.008	0.001 **	-0.023	0.003 **	-0.051	0.006 **
Q1	Donut	0.036	0.001 **	0.006	0.006	0.023	0.008 **
Q2	ICR	-0.009	0.002 **	-0.022	0.003 **	-0.041	0.007 **
Q2	Donut	0.062	0.002 **	0.018	0.007 **	0.026	0.008 **
Q3	ICR	-0.016	0.002 **	-0.019	0.003 **	-0.043	0.007 **
Q3	Donut	0.077	0.002 **	0.002	0.007	0.008	0.009
Q4	ICR	-0.025	0.002 **	-0.019	0.003 **	-0.051	0.009 **
Q4	Donut	0.074	0.002 **	-0.015	0.007 *	-0.011	0.011
Q1	ICR	0.030	0.003 **	-0.011	0.005 *	-0.034	0.010 **
Q1	Donut	-0.028	0.003 **	-0.042	0.010 **	-0.053	0.014 **
Q2	ICR	0.010	0.003 **	-0.004	0.005	-0.023	0.011 *
Q2	Donut	-0.048	0.003 **	-0.056	0.011 **	-0.073	0.015 **
Q3	ICR	-0.005	0.004	-0.003	0.006	-0.023	0.012
Q3	Donut	-0.092	0.003 **	-0.065	0.013 **	-0.086	0.018 **
Q4	ICR	0.006	0.004	0.011	0.007	-0.008	0.014
Q4	Donut	-0.155	0.004 **	-0.124	0.015 **	-0.164	0.025 **

Notes: Results of quarterly OLS, reduced form, and instrumental variables regressions of log consumption change on log change in ICR and donut prices (or, in the reduced form case, the log change in ICR and donut IV prices), separately for low-spending individuals and high-spending individuals. Individuals with positive consumption in each quarter only. All years pooled. N=919,650 for low-spending group; N=294,898 for high-spending group. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

The patterns in the IV results generally track the patterns in the reduced form results across spending group, coverage phase (ICR vs. donut), and quarter combinations. E.g., in both the reduced form and IV results, low-spenders' response to the ICR price change is negative and significant, but flat over time, while their response to the donut hole price change is small and positive. Similarly, the reduced form and IV results show that high-spenders'

response to the ICR price change is low and decreasing in magnitude (and, conversely, their response to the donut price change is steeply increasing in magnitude) over the course of the year, as enrollees exit the ICR and enter the donut.

Both the reduced form and IV results deviate somewhat from the OLS results, which are often positive, contrary to expectations about downward-sloping demand curves (a significant exception is high-spenders' response to the donut price change). This discrepancy between the OLS and IV results is likely generated by the fact that, within each year pair over which differences in prices and consumption are taken, the instrumental variables prices for both years hold consumption weights and retail prices fixed at the first year's levels. The OLS approach, in not doing so, permits an endogeneity problem wherein, for example, enrollees experiencing upward shocks in health care needs require more expensive drugs, inducing a positive association between the year-to-year consumption change and price change.

C Categorical Donut Hole Coverage

As noted in the text, some of the donut hole coverage variation observed in our sample pertains to broad categories of drugs and would be easily understood by enrollees, but the sample also includes many changes that would be difficult for enrollees to translate into prices.

In 2006-7, all donut hole coverage changes are of the former, “stark” variety, in that they entail plans adding or dropping an entire category or more of drugs to the donut hole coverage; for example, consider Appendix Table 11a, which shows the count of sample enrollees in 2006-7 by 2006 gap coverage (across rows) and 2007 gap coverage (across columns). The Table shows that 35,325 individuals were in plans with no donut hole coverage in 2006 and generic donut hole coverage in 2007 – adding generic coverage implies large average price decreases for the donut hole of about \$10 per 30-day supply. These decreases would be simple for individuals to understand given knowledge of the prices they face for branded and generic drugs in the ICR and an understanding of which drugs are generic. In contrast, some plans in 2007-8 and even more in 2008-9 had slight alterations in their coverage in the donut hole which did not entail large average price changes and which were not generally easily understandable; see the count of individuals according to year 1 and year 2 gap coverage in Appendix Table 11b and Appendix Table 11c. For example, 26,551 enrollees changed from “All Generic” coverage in 2008 to “Many Generic” coverage in 2009. This did not serve to universally increase average prices in the donut hole – some plans still decreased copays for covered generics while removing coverage for others – and the average price increase across plans was small, around \$0.56. Further, this type of coverage change would require a more complicated calculation for individuals to respond to it than a stark coverage change such as removing/adding coverage for an entire class of drugs – it is arguably surely easier for enrollees to identify which drugs are branded and generic than to identify which generic drugs are “Many Generic” or which branded drugs are “Few Brand.”⁴¹

Appendix Table 11a: Enrollment by Donut Hole Coverage Changes, 2006-7

		2007 Gap Coverage, 2006-7 Sample			
		Generics,			
		Brand & Gen	Pref. Brands	Generic	None
2006 Gap Coverage, 2006-7 Sample	Brand & Gen	7,832	18	12,895	6,773
	Generic	152	132	31,172	3,093
	None	1,851	116	35,325	352,273

Notes: Count of sample enrollees with given gap coverage in 2006 and 2007 chosen plan(s). 2006 gap coverage designated by row value; 2007 gap coverage designated by column value.

⁴¹The fact that essentially none of the donut coverage variation in 2008-9 is of this stark form (in contrast to 2006-7 and 2007-8) may account for the observation in Table 5 that end of year marginal price responses are somewhat smaller at the end of the year for the 2008-9 sample than for the 2006-7 and 2007-8 samples.

Appendix Table 11b: Enrollment by Donut Hole Coverage Changes, 2007-8

		2008 Gap Coverage, 2007-8 Sample					
2007 Gap Coverage, 2007-8 Sample		All Generics and Some Brands			All Preferred Generics	All Generics	None
		Some Generics					
	Brand & Gen	89	1	666	7,556	7,746	
	Generics, Pref. Brands	1	0	392	7	136	
	Generic	27,909	19	62,653	52,386	24,629	
	None	218	9	2,690	5,211	934,364	

Notes: Count of sample enrollees with given gap coverage in 2007 and 2008 chosen plan(s). 2007 gap coverage designated by row value; 2008 gap coverage designated by column value.

Appendix Table 11c: Enrollment by Donut Hole Coverage Changes, 2008-9

		2009 Gap Coverage, 2008-9 Sample					
		All Generics and Few Brands			Many Generics and Few Brands		
		Some Generics	All Generics		Many Generics		None
2008 Gap Coverage, 2008-9 Sample	Some Generics	985	48	0	25,710	0	1,651
	All Preferred Generics	136	252	14	55,588	384	6,336
	All Generics	701	30,506	44	26,551	0	7,230
	All Generics and Some Brands	28	0	0	0	0	13
	None	3,099	775	55	2,452	2	966,640

Notes: Count of sample enrollees with given gap coverage in 2008 and 2009 chosen plan(s). 2008 gap coverage designated by row value; 2009 gap coverage designated by column value.

In incorporating donut coverage variation in our sample into our regression analysis, we focus on the sort of “stark” variation observed primarily in 2006-7 and 2007-8, as it maps more clearly into price changes and is more likely to be understood by enrollees. Sensitivity analysis where we include separate coverage change variables for “stark” coverage changes and for “non-stark” coverage changes leave the included variables unchanged – the “non-stark” variables capturing subtle donut hole price changes are not statistically significant.

D Taylor Expansion in Structural Model

Letting \mathbf{z}_y be the vector of year y prices, we linearize the specification around \mathbf{z}_1 . We use the following Taylor expansion:

$$\begin{aligned}
q_{it,2}(\mathbf{z}_2) &= f(\mathbf{z}_2) = q_{it,1} * \exp\left(a_t + X_{it} * \delta_t + b_{it,ICR} * \left(\frac{\log(P_{i,ICR,2})}{\log(P_{i,ICR,1})}\right)\right) \\
&* q_{it,1} * \exp\left(b_{it,Donut} * \left(\frac{\log(P_{i,Donut,2})}{\log(P_{i,Donut,1})}\right)\right) \\
&* q_{it,1} * \exp(\theta_{Ded,t} * dedchg_i + \theta_{stark,t} * stark_i) \\
&* q_{it,1} * \exp(u_{it,2}) \\
&\cong f(\mathbf{z}_1) + \left(\frac{\partial f(\mathbf{z}_1)}{\partial P_{ICR}}\right) (P_{i,ICR,2} - P_{i,ICR,1}) + \left(\frac{\partial f(\mathbf{z}_1)}{\partial P_{Don}}\right) (P_{i,Don,2} - P_{i,Don,1}) \\
&+ \left(\frac{\partial f(\mathbf{z}_1)}{\partial dedchg}\right) dedchg_i + \left(\frac{\partial f(\mathbf{z}_1)}{\partial stark}\right) stark_i.
\end{aligned}$$

The Taylor expansion yields the following form for year 2 consumption:

$$\begin{aligned}
q_{it,2}(\mathbf{z}_2) &\cong q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) \\
&+ q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * \left(\frac{P_{i,ICR,2} - P_{i,ICR,1}}{P_{i,ICR,1}}\right) * b_{it,ICR} \\
&+ q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * \left(\frac{P_{i,Don,2} - P_{i,Don,1}}{P_{i,Don,1}}\right) * b_{it,Don} \\
&+ q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * (dedchg_i * \theta_{Ded,t} + stark_i * \theta_{Stark,t}) \\
&= q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) (1 - b_{it,ICR} - b_{it,Don}) \\
&+ \left(\frac{q_{i,1}}{P_{i,ICR,1}}\right) * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * P_{i,ICR,2} * b_{it,ICR} \\
&+ \left(\frac{q_{i,1}}{P_{i,Don,1}}\right) * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * P_{i,Don,2} * b_{it,Don} \\
&+ q_{it,1} * \exp(a_t + X_{it} * \delta_t + u_{it,2}) * (dedchg_i * \theta_{Ded,t} + stark_i * \theta_{Stark,t})
\end{aligned}$$

E Robustness

In this Section, we explore the sensitivity of our methodology to the assumptions that enrollees in our estimation sample have no uncertainty about marginal coverage phase and that current and marginal coverage phase assignments are fixed and exogenous with respect to out-of-pocket prices.

E.1 Limited Uncertainty about Marginal Coverage Phase

In the regression estimations used to obtain our structural parameter estimates, we noted that our restriction to “low” and “high” spending individuals does not guarantee that those individuals have no uncertainty about marginal coverage phase. “Low” spending individuals cross the donut threshold in year 2 3% of the time; “high” spending individuals do *not* cross the donut threshold in year 2 14% of the time. In this Section, we investigate sensitivity to classification of the low and high-spending individuals to determine whether uncertainty or switching behavior is likely to bias our results. In order to perform this check, we first regress an indicator for ending the year in the “appropriate” (ICR for low-spending individuals, donut for high-spending individuals) coverage phase on all demographic and spending controls. We reduce the “error”, as defined by the individual ending the year on the wrong side of the donut threshold, by 50% and by 25% by restricting the sample based on the predicted probability of ending in the appropriate marginal coverage phase, and display the results of the quarterly regression for these samples.

The results are shown in Appendix Table 12. Of the 48 coefficients estimated, two are statistically significantly different than in the baseline sample – the deductible change response is larger in magnitude in the restricted samples, significantly so in the 25% restricted sample for Q2, and the Q2 ICR response by low-spending individuals is significantly smaller in the 50% restricted sample. Otherwise, the point estimates are similar in magnitude and exhibit similar dynamic patterns – the marginal price coefficients in the 25% restricted sample are slightly smaller in magnitude, while the marginal price coefficients in the 50% restricted sample are sometimes smaller, sometimes larger, than in the baseline sample.

Appendix Table 12: Results of Quarterly ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – High- and Low-Spending Enrollees, Pooled Years, Baseline and Restricted Samples

Period	Price	Baseline Thresholds [N=1,214,548]				Pred (Error Reduce 25%) [N=988,123]				Pred (Error Reduce 50%) [N=700,114]			
		Low-Spending		High-Spending		Low-Spending		High-Spending		Low-Spending		High-Spending	
		Enrollee Response		Enrollee Response		Enrollee Response		Enrollee Response		Enrollee Response		Enrollee Response	
		Coef	SE	Coef	SE	Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	-0.056	0.005 **	-0.022	0.007 **	-0.056	0.006 **	-0.019	0.009 *	-0.049	0.007 **	-0.023	0.013
Q1	Donut	0.030	0.007 **	-0.040	0.009 **	0.027	0.007 **	-0.032	0.011 **	0.021	0.009 *	-0.031	0.017
Q1	Ded. Chg	-0.035	0.005 **	-0.035	0.005 **	-0.046	0.006 **	-0.046	0.006 **	-0.047	0.007 **	-0.047	0.007 **
Q1	Stark	-0.015	0.006 *	-0.015	0.006 *	-0.026	0.012 *	-0.026	0.012 *	-0.009	0.015	-0.009	0.015
Q2	ICR	-0.051	0.005 **	-0.007	0.007	-0.044	0.006 **	-0.002	0.009	-0.029	0.007 **	0.011	0.015
Q2	Donut	0.029	0.007 **	-0.051	0.015 **	0.030	0.008 **	-0.029	0.016	0.025	0.009 **	-0.029	0.023
Q2	Ded. Chg	-0.009	0.005	-0.009	0.005	-0.027	0.006 **	-0.027	0.006 **	-0.014	0.007 *	-0.014	0.007 *
Q2	Stark	0.001	0.008	0.001	0.008	0.010	0.011	0.010	0.011	0.011	0.015	0.011	0.015
Q3	ICR	-0.048	0.005 **	0.002	0.008	-0.038	0.006 **	0.016	0.010	-0.036	0.007 **	0.008	0.016
Q3	Donut	0.014	0.007	-0.091	0.014 **	0.011	0.008	-0.075	0.015 **	0.001	0.010	-0.085	0.020 **
Q3	Ded. Chg	-0.008	0.005	-0.008	0.005	-0.016	0.006 **	-0.016	0.006 **	-0.018	0.008 *	-0.018	0.008 *
Q3	Stark	0.000	0.007	0.000	0.007	-0.005	0.011	-0.005	0.011	-0.014	0.014	-0.014	0.014
Q4	ICR	-0.062	0.007 **	0.032	0.010 **	-0.054	0.008 *	0.050	0.013 *	-0.052	0.009 **	0.050	0.019 **
Q4	Donut	0.006	0.010	-0.183	0.017 **	0.007	0.011	-0.172	0.026 **	0.007	0.012	-0.184	0.039 **
Q4	Ded. Chg	-0.004	0.005	-0.004	0.005	-0.015	0.007 *	-0.015	0.007 *	-0.011	0.008	-0.011	0.008
Q4	Stark	-0.041	0.009 **	-0.041	0.009 **	-0.058	0.014 **	-0.058	0.014 **	-0.071	0.017 **	-0.071	0.017 **

Notes: Results of quarterly regressions of log consumption change on log change in ICR and donut prices, as well as changes in stark gap coverage and deductible coverage, main regression estimation sample ("Baseline") vs. more restricted samples. The sample denoted "Error Reduce 25%" uses the results of a regression of an indicator for ending the year in the appropriate coverage phase on year 1 demographic and spending variables to reduce the "Error" rate (defined as ending the year on the wrong side of the donut threshold) by X%. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

On balance, the results of this specification check indicate that our results are not sensitive to further restriction of our sample to reduce marginal coverage phase switching behavior.

E.2 Allowing Current and Marginal Coverage Phase to Respond to Out-of-Pocket Prices

Recall the derivation of the dynamic structural model in Section 4.1. The solution for optimal consumption shown in equation (2) assumed that individuals end each period in the interior of a coverage phase (i.e., no bunching), and also that current and marginal coverage phases do not respond to out-of-pocket prices (i.e., no switching). The restriction of the estimation sample to individuals ending year 1 well away from the donut threshold is intended to limit violations of these assumptions, and the evidence in Section E.1 above provides reassurance on this point. In this Section, we estimate a richer model that allows for both bunching and switching, and find ultimately that our results are unaffected.

Consider a more general version of the model in Section 4.1, where consumers choose consumption in each period according to the following value functions:

$$\begin{aligned}
 V_T(X_T, A_T) &= W_T(X_T, A_T) = \max_q u(q, A_T) - E^{OOP}(X_T, q) \\
 V_t(X_t, A_t) &= \max_q u(q, A_t) - E^{OOP}(X_t, q) + \beta \int W_{t+1}(X_t + q * R, A_{t+1}) dF(A_{t+1}) \quad \forall t < T \\
 W_t(X_t, A_t) &= \max_q u(q, A_t) - E^{OOP}(X_t, q) + \int W_{t+1}(X_t + q * R, A_{t+1}) dF(A_{t+1}).
 \end{aligned}$$

Both V_t and W_t are continuous everywhere and differentiable everywhere except $X_t = \bar{X}$. A_t denotes a set of parameters impacting the utility function in period t (e.g., A_t could scale

marginal utility in period t). Denote $u_t(q) = u(q, A_t)$. We allow for uncertainty regarding prescription drug needs in future periods using this term. For the sake of exposition, we begin by assuming that there is a single coverage phase kink, a convex kink at \bar{X} where the out-of-pocket price changes from p_1 to $p_2 > p_1$. If an individual has spent X_t on drugs up until period t and purchases q_t units in period t , her period t out-of-pocket expenditure will be:

$$E^{OOP}(X_t, q_t) = \begin{cases} p_1 * q_t & \text{if } X_t + R * q_t \leq \bar{X} \\ p_1 * \frac{\bar{X} - X_t}{R} + p_2 * \left(q_t - \frac{\bar{X} - X_t}{R} \right) & \text{if } X_t + R * q_t > \bar{X} \text{ and } X_t \leq \bar{X} \\ p_2 * q_t & \text{if } X_t > \bar{X}. \end{cases}$$

We will then generalize the solution to allow for deductible and catastrophic coverage phases in addition to the ICR and donut.

Consider first period T ; at the beginning of period T , all uncertainty regarding prescription drug needs for the year will have been resolved. If $X_T \geq \bar{X}$, then the consumer faces a linear price of p_2 per unit and optimal consumption will equal $q_T^* = u_T'^{-1}(p_2)$. If $X_T < \bar{X}$, then the solution will be the piecewise nonlinear budget set solution as in equation (1):

$$q_T^*(X_T) = \begin{cases} u_T'^{-1}(p_1) & \text{if } u_T' \left(\frac{\bar{X} - X_T}{R} \right) \leq p_1 \text{ and } X_T \leq \bar{X} \\ \frac{\bar{X} - X_T}{R} & \text{if } p_1 < u_T' \left(\frac{\bar{X} - X_T}{R} \right) \leq p_2 \text{ and } X_T \leq \bar{X} \\ u_T'^{-1}(p_2) & \text{if } u_T' \left(\frac{\bar{X} - X_T}{R} \right) > p_2 \text{ or } X_T > \bar{X}. \end{cases}$$

Note that there is some probability of bunching in the final period if $X_T < \bar{X}$. Denote $\tilde{q}_t^c = u_t'^{-1}(p_c)$.

Next consider period $T - 1$. If $X_{T-1} \geq \bar{X}$, then the consumer faces a linear price p_2 in the remaining period and $W_T'(X, A_T) = 0$ for all A_T . The solution will be $q_{T-1}^* = \tilde{q}_{T-1}^2$. If $X_{T-1} < \bar{X}$, then we must consider three cases. If the individual consumes past the coverage threshold in period $T - 1$, then we again have $W_T'(X, A_T) = 0$ for all A_T and the solution will be $q_{T-1}^* = \tilde{q}_{T-1}^2$. If the individual bunches in period $T - 1$, then mechanically $q_{T-1}^* = (\bar{X} - X_{T-1})/R$. Finally, if the individual remains in the ICR in period $T - 1$, then we must in turn consider three possibilities for period T consumption behavior – the case where the individual remains in the ICR in period T as well: $\{stay_T\}$; the case where the individual bunches in period T : $\{bunch_T\}$; and the case where the individual crosses the threshold in period T : $\{cross_T\}$. The objective function conditional on remaining in the ICR in period $T - 1$ is:

$$\begin{aligned} & \max_q u_{T-1}(q) - p_1 * q \\ & + \beta * Pr\{stay_T\} * \mathbb{E}(W_T(X_{T-1}, A_T) | stay_T) \\ & + \beta * Pr\{bunch_T\} * \mathbb{E}(W_T(X_{T-1}, A_T) | bunch_T) \\ & + \beta * Pr\{cross_T\} * \mathbb{E}(W_T(X_{T-1}, A_T) | cross_T). \end{aligned}$$

Substituting in the results from above, we have:

$$\begin{aligned}
& \max_q u_{T-1}(q) - p_1 * q \\
& + \Pr\{stay_T\} * \beta * \int (u_T(\tilde{q}_T^1) - p_1 * \tilde{q}_T^1) dF(A_T) \\
& + \Pr\{bunch_T\} * \beta * \int \left(u_T \left(\frac{\bar{X} - X_{T-1}}{R} - q \right) - p_1 * \left(\frac{\bar{X} - X_{T-1}}{R} - q \right) \right) dF(A_T) \\
& + \Pr\{cross_T\} * \beta * \int \left(u_T(\tilde{q}_T^2) - p_1 * \left(\frac{\bar{X} - X_{T-1}}{R} - q \right) - p_2 * \left(\tilde{q}_T^2 - \frac{\bar{X} - X_{T-1}}{R} + q \right) \right) dF(A_T).
\end{aligned}$$

By the envelope theorem, this implies the following first-order condition:

$$\begin{aligned}
u'_{T-1}(q) &= p_1 + \Pr\{bunch_T\} * \beta * \int \left(u'_T \left(\frac{\bar{X} - X_{T-1}}{R} - q \right) - p_1 \right) dF(A_T) \\
&+ \Pr\{cross_T\} * \beta * (p_2 - p_1) \\
&= (1 - \beta * \Pr\{cross_T\}) * p_1 + \Pr\{cross_T\} * \beta * p_2 \\
&+ \beta * \Pr\{bunch_T\} * \mathbb{E} \left(u'_T \left(\frac{\bar{X} - X_{T-1}}{R} - q \right) - p_1 \right).
\end{aligned}$$

Let $q_{T-1}^S(X_{T-1})$ be the unique solution to the above equation. To determine the region of A_{T-1} for which the individual will remain in the ICR ($stay_{T-1}$), bunch at the donut kink ($bunch_{T-1}$), or cross into the donut hole ($cross_{T-1}$), we consider the limits of the first-order condition as q approaches $(\bar{X} - X_{T-1})/R$ from the left-hand side and right-hand side. The limit of the period $T - 1$ first-order condition above as we approach the kink from the left-hand side is:

$$u'_{T-1}(q) - (1 - \beta) * p_1 - \beta * p_2,$$

as the probability of entering the donut in period T goes to one (and, accordingly, as $\Pr\{bunch_T\}$ goes to zero). The limit of the period $T - 1$ first-order condition as we approach the kink from the right-hand side is simply:

$$u'_{T-1}(q) - p_2$$

since $W'_T = 0$ for all q such that $X_{T-1} + R * q > \bar{X}$. Then optimal consumption in period $T - 1$ will be:

$$q_{T-1}^*(X_{T-1}) = \begin{cases} q_{T-1}^S(X_{T-1}) & \text{if } u'_{T-1} \left(\frac{\bar{X} - X_{T-1}}{R} \right) \leq (1 - \beta) * p_1 + \beta * p_2 \text{ and } X_{T-1} \leq \bar{X} \\ \frac{\bar{X} - X_{T-1}}{R} & \text{if } (1 - \beta) * p_1 + \beta * p_2 < u'_T \left(\frac{\bar{X} - X_{T-1}}{R} \right) \leq p_2 \text{ and } X_{T-1} \leq \bar{X} \\ \tilde{q}_{T-1}^2 & \text{if } u'_T \left(\frac{\bar{X} - X_{T-1}}{R} \right) > p_2 \text{ or } X_{T-1} > \bar{X} \end{cases}.$$

This analysis yields the following general solution.

Theorem 1 For any period $t < T$, the optimal consumption path will be:

$$q_t^*(X_t) = \begin{cases} q_t^S(X_t) & \text{if } X_t \leq \bar{X} \text{ and } u'_t\left(\frac{\bar{X}-X_t}{R}\right) \leq (1-\beta)*p_1 + \beta*p_2 \\ \frac{\bar{X}-X_t}{R} & \text{if } X_t \leq \bar{X} \text{ and } (1-\beta)*p_1 + \beta*p_2 < u'_t\left(\frac{\bar{X}-X_t}{R}\right) \leq p_2 \\ \tilde{q}_t^2 & \text{if } X_t > \bar{X} \text{ or } u'_t\left(\frac{\bar{X}-X_t}{R}\right) > p_2 \end{cases}.$$

with $q_t^S(X_t)$ defined by

$$\begin{aligned} u'_t(q_t^S(X_t)) &= \left(1 - \beta * \sum_{i=t+1}^T \Pr\{\text{cross}_i|t\}\right) * p_1 + \beta * \sum_{i=t+1}^T \Pr\{\text{cross}_i|t\} * p_2 \\ &\quad + \beta * \sum_{i=t+1}^T \Pr\{\text{bunch}_i|t\} * \mathbb{E}\left(u'_i\left(\frac{\bar{X}-X_{i-1}}{R} - q_{i-1}\right) - p_1|t\right). \end{aligned}$$

Proof 1 We prove this result by induction. Suppose that the result holds in $t+1$, and consider the optimal consumption for period t . If $X_t \geq \bar{X}$, then as usual the consumer faces a linear price p_2 in all remaining periods and $W'_{t+1}(X, A_{t+1}) = 0$ for all A_{t+1} ; the solution will be $q_t^* = u'^{-1}_t(p_2) = \tilde{q}_t^2$. If $X_t < \bar{X}$, then we must consider three cases. If the individual consumes past the coverage threshold in period t , then we again have $W'_{t+1}(X, A_{t+1}) = 0$ for all A_{t+1} and the solution will be $q_t^* = \tilde{q}_t^2$. If the individual bunches in period t , then mechanically $q_t^* = (\bar{X} - X_t)/R$. If, however, the individual remains in the ICR in period t , then we must consider the three potential outcomes in period $t+1$: $\{\text{stay}_{t+1}\}$, $\{\text{bunch}_{t+1}\}$, and $\{\text{cross}_{t+1}\}$. The objective function is:

$$\begin{aligned} &\max_q u_t(q) - p_1 * q \\ &+ \beta * \Pr\{\text{stay}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{stay}_{t+1}) \\ &+ \beta * \Pr\{\text{bunch}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{bunch}_{t+1}) \\ &+ \beta * \Pr\{\text{cross}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{cross}_{t+1}). \end{aligned}$$

Again applying the envelope theorem, the first-order condition for this objective function is:

$$\begin{aligned} u'_t(q) &= p_1 \\ &- \beta * \Pr\{\text{stay}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{stay}_{t+1}) \\ &- \beta * \Pr\{\text{bunch}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{bunch}_{t+1}) \\ &- \beta * \Pr\{\text{cross}_{t+1}|t\} * \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{cross}_{t+1}). \end{aligned} \tag{9}$$

If the individual either crosses or bunches in period $t+1$, then $\mathbb{E}(W'_i(X_i, A_i)|t) = 0$ for all $i > t+1$. Accordingly, if $\{\text{cross}_{t+1}\} = 1$, then

$$\mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{cross}_{t+1}) = -(p_2 - p_1) \tag{10}$$

and if $\{\text{bunch}_{t+1}\} = 1$, then

$$\mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1})|\text{bunch}_{t+1}) = -\mathbb{E}\left(u'_{t+1}\left(\frac{\bar{X} - X_{t+1}}{R}\right)|t\right) + p_1. \tag{11}$$

Finally, if $\{stay_{t+1}\} = 1$, then by the inductive hypothesis we have

$$\begin{aligned} & \mathbb{E}(W'_{t+1}(X_t + R * q, A_{t+1}) | t + 1, stay_{t+1}) \\ = & \mathbb{E} \left(\sum_{i=t+2}^T Pr\{bunch_i | t + 1, stay_{t+1}\} * \mathbb{E} \left(-u'_i \left(\frac{\bar{X} - X_i}{R} \right) + p_1 | t + 1, stay_{t+1} \right) | t \right) \quad (12) \\ & + \mathbb{E} \left(\sum_{i=t+2}^T Pr\{cross_i | t + 1, stay_{t+1}\} * (p_1 - p_2) | t \right). \end{aligned}$$

Substituting equations (10), (11), and (12) into equation (9), and applying the law of iterated expectations, we obtain

$$\begin{aligned} u'_t(q) &= p_1 \\ &- \beta * \mathbb{E} \left(\sum_{i=t+2}^T Pr\{bunch_i\} * \left(-u'_i \left(\frac{\bar{X} - X_i}{R} \right) + p_1 \right) | t \right) \\ &- \beta * \mathbb{E} \left(\sum_{i=t+2}^T Pr\{cross_i\} * (p_1 - p_2) | t \right) \quad (13) \\ &+ \beta * Pr\{bunch_{t+1} | t\} * \left(\mathbb{E} \left(u'_{t+1} \left(\frac{\bar{X} - X_{t+1}}{R} \right) | t \right) - p_1 \right) \\ &+ \beta * Pr\{cross_{t+1} | t\} * (p_2 - p_1). \end{aligned}$$

The second and third lines follow from the fact that $Pr\{stay_{t+1} | t\} * Pr\{bunch_i | stay_{t+1}, t\} = Pr\{stay_{t+1} \cap bunch_i | t\} = Pr\{stay_{t+1} | bunch_i, t\} * Pr\{bunch_i | t\} = Pr\{bunch_i | t\}$ and similarly $Pr\{stay_{t+1} | t\} * Pr\{cross_i | stay_{t+1}, t\} = Pr\{cross_i | t\}$ for all $i > t + 1$ (by definition, one cannot cross or bunch in any period unless they have “stayed” in all previous periods). This expression simplifies to

$$\begin{aligned} u'_t(q) &= p_1 \\ &+ \beta * \mathbb{E} \left(\sum_{i=t+1}^T Pr\{bunch_i | t\} * \left(u'_i \left(\frac{\bar{X} - X_i}{R} \right) - p_1 \right) | t \right) \\ &+ \beta * \sum_{i=t+1}^T Pr\{cross_i | t\} * (p_2 - p_1). \end{aligned}$$

This implies that the optimal consumption level if the individual remains in the ICR in period t will be $q_t^S(X_t)$ such that

$$\begin{aligned} u'_t(q_t^S(X_t)) &= \left(1 - \beta * \sum_{i=t+1}^T Pr\{cross_i | t\} \right) * p_1 + \beta * \sum_{i=t+1}^T Pr\{cross_i | t\} * p_2 \\ &+ \beta * \sum_{i=t+1}^T Pr\{bunch_i | t\} * \mathbb{E} \left(u'_i \left(\frac{\bar{X} - X_{i-1}}{R} - q_{i-1} \right) - p_1 | t \right) \end{aligned}$$

which, in turn, implies the full solution⁴²

$$q_t^*(X_t) = \begin{cases} q_t^S(X_t) & \text{if } X_t \leq \bar{X} \text{ and } u'_t \left(\frac{\bar{X} - X_t}{R} \right) \leq (1 - \beta) * p_1 + \beta * p_2 \\ \frac{\bar{X} - X_t}{R} & \text{if } X_t \leq \bar{X} \text{ and } (1 - \beta) * p_1 + \beta * p_2 < u'_t \left(\frac{\bar{X} - X_t}{R} \right) \leq p_2 \\ \tilde{q}_t^2 & \text{if } X_t > \bar{X} \text{ or } u'_t \left(\frac{\bar{X} - X_t}{R} \right) > p_2 \end{cases}.$$

⁴²We apply the same reasoning as in the above analysis for $T - 1$ to define the ranges of $u'_t(\cdot)$ such that the individual will choose to stay, bunch, or cross in period t .

Thus completing the proof.

We now describe how we generalize this model to allow for the deductible and catastrophic coverage phases, estimate the richer model on our sample of low- and high-spending individuals, and use the resulting parameters to simulate consumption for the full sample (including individuals near the donut hole kink) as in Section 6. While our model in principle can allow for bunching, in practice the probability of bunching in the sample of low- and high-spending individuals is so low ($< 0.1\%$ according any model of expectations) that we omit this term.⁴³ We then have

$$q_t^*(X_t) = \begin{cases} q_t^S(X_t) & \text{if } X_t \leq \bar{X} \text{ and } u'_t\left(\frac{\bar{X}-X_t}{R}\right) \leq (1-\beta) * p_1 + \beta * p_2 \\ \tilde{q}_t^2 & \text{if } X_t > \bar{X} \text{ or } u'_t\left(\frac{\bar{X}-X_t}{R}\right) > p_2 \end{cases}$$

with $q_t^S(X_t)$ such that

$$u'_t(q_t^S(X_t)) = \left(1 - \beta * \sum_{i=t+1}^T \Pr\{cross_i|t\}\right) * p_1 + \beta * \sum_{i=t+1}^T \Pr\{cross_i|t\} * p_2.$$

It is useful to rewrite this solution such that marginal utility is equal to the appropriate virtual price given the coverage phase in which the individual ends the current period and the year:

$$u'(q_t^*) = \Pr\{ICR_t\} [(1 - \beta * \Pr\{donut_T|ICR_t\}) * p_1 + \beta * \Pr\{donut_T|ICR_t\} * p_2] + \Pr\{donut_t\} * p_2$$

where, in this notation, $\Pr\{c_t\}$ is the probability that the individual ends period t in coverage phase c . Again, all bunching terms are omitted. The deductible and catastrophic phases are conceptually straightforward to incorporate into our model because they are convex kinks and thus do not permit bunching. Repeating the above model derivation allowing for these additional coverage phases yields:

$$\begin{aligned} u'(q_t^*) = & \Pr\{ded_t\} * (1 - \beta * (\Pr\{ICR_T|ded_t\} + \Pr\{don_T|ded_t\} + \Pr\{cat_T|ded_t\})) * p_0 \\ & + \Pr\{ded_t\} * \beta * [\Pr\{ICR_T|ded_t\} * p_1 + \Pr\{don_T|ded_t\} * p_2 + \Pr\{cat_T|ded_t\} * p_3] \\ & + \Pr\{ICR_t\} * (1 - \beta * (\Pr\{don_T|ICR_t\} + \Pr\{cat_T|ICR_t\})) * p_1 \\ & + \Pr\{ICR_t\} * \beta * [\Pr\{don_T|ICR_t\} * p_2 + \Pr\{cat_T|ICR_t\} * p_3] \\ & + \Pr\{don_t\} * ((1 - \beta * \Pr\{cat_T|don_t\}) * p_2 + \beta * \Pr\{cat_T|don_t\} * p_3) \\ & + \Pr\{cat_t\} * p_3. \end{aligned}$$

This expression simplifies to our result from equation (2): $q^* = u'^{-1}((1-\beta) * CP + \beta * MP)$. However, it provides greater detail on where each individual price enters the optimal consumption function. Suppose, as in our structural modeling exercise above, that $u(q)$ is

⁴³The estimation strategy below can be modified to allow for bunching by estimating an auxiliary model of bunching as a function of prices.

quadratic in q . We can then rearrange and differentiate to obtain:

$$\begin{aligned}
\frac{\partial q_t^*}{\partial p_1} = & \eta * \frac{\partial Pr\{ded_t\} * (1 - \beta * (Pr\{ICR_T|ded_t\} + Pr\{don_T|ded_t\} + Pr\{cat_T|ded_t\}))}{\partial p_1} * p_0 \\
& + \eta * \beta * \left[\frac{\partial Pr\{ded_t \cap ICR_T\}}{\partial p_1} * p_1 + \frac{\partial Pr\{ded_t \cap don_T\}}{\partial p_1} * p_2 + \frac{\partial Pr\{ded_t \cap cat_T\}}{\partial p_1} * p_3 \right] \\
& + \eta * \left[\frac{\partial Pr\{ICR_t\}}{\partial p_1} - \beta * \left(\frac{\partial Pr\{ICR_t \cap don_T\}}{\partial p_1} + \frac{\partial Pr\{ICR_t \cap cat_T\}}{\partial p_1} \right) \right] * p_1 \\
& + \eta * \beta * \left[\frac{\partial Pr\{ICR_t \cap don_T\}}{\partial p_1} * p_2 + \frac{\partial Pr\{ICR_t \cap cat_T\}}{\partial p_1} * p_3 \right] \\
& + \eta * \left(\frac{\partial Pr\{don_t\}}{\partial p_1} - \beta * \frac{\partial Pr\{don_t \cap cat_T\}}{\partial p_1} \right) * p_2 + \eta * \beta * \frac{\partial Pr\{don_t \cap cat_T\}}{\partial p_1} * p_3 \\
& + \eta * \frac{\partial Pr\{cat_t\}}{\partial p_1} * p_3 \\
& + \eta * \beta * Pr\{ded_t \cap ICR_T\} \\
& + \eta * Pr\{ICR_t\} * (1 - \beta * (Pr\{don_T|ICR_t\} + Pr\{cat_T|ICR_t\}))
\end{aligned} \tag{14}$$

and similarly for $\partial q_t^*/\partial p_2$. The key distinction between this specification and the one estimated in Section 6 is that the current, richer specification allows for endogenous coverage phase switching, as each coverage phase probability term is allowed to respond endogenously to coverage phase prices.

In order to estimate this richer model, we use the following procedure: (1) we estimate how each coverage phase probability term (e.g., $Pr\{ICR_T|ded_t\}$, $Pr\{don_t\}$, etc.) responds to p_1 and p_2 . We use the same exact regression specification and controls as in Section 4, but use the relevant probability term as the left-hand-side variable. The probability term is calculated assuming perfect foresight as our model of expectations, as in Section 6. Next, we (2) perform the GMM estimation as in Section 6, but instead of using equations (4) and (5) in the GMM objective function directly, we replace the left-hand sides of equations (4) and (5) with equation (14) and its analog for the donut price. We use this richer specification as the GMM objective function to recover η and β . Finally (3), using the new estimates of η and β , we simulate consumption for all sample individuals and compare to observed consumption.⁴⁴

The results of this procedure are overall quite similar to the results when we do not allow for endogenous coverage phase switching. The estimates of η and β are not statistically significantly different from the baseline estimates in Section 6 (the richer model yields $\beta = 0.313$ and $\bar{\eta} = -1.77$, whereas the simpler model estimates were $\beta = 0.312$ (SE=0.08) and $\bar{\eta} = -1.66$ (SE=0.15)). The comparison of actual to simulated spending are shown, for each model, in Appendix Table 13. If anything, the simulation of consumption both in- and outside the regression sample performs slightly *worse* than the simpler approach – the mean squared error with the richer specification is 1% higher in the regression sample, and 5% higher in the sample of individuals near the donut threshold.

⁴⁴Note that, whether we use the simple GMM procedure in Section 6 or the richer GMM procedure outlined here to recover η and β , the simulations of consumption in- and outside the regression sample always allow for bunching and switching behavior.

Appendix Table 13: Comparison of Actual to Predicted Spending – Basic Structural Model and Richer Specification

	Baseline Structural Model Comparison				Richer Structural Model Comparison			
	Actual	Predicted			Actual	Predicted		
		Mean	% diff	MSE		Mean	% diff	MSE
In-Sample (Low/High Spending Enrollees)	1,718.239	1,703.561	-0.85%	996,251	1,718.239	1,694.167	-1.40%	1,006,810
Out-of- Sample (Medium Spending Enrollees)	2,324.124	2,328.278	0.18%	990,186	2,324.124	2,316.333	-0.34%	1,035,317
All Spending Groups	1,972.520	1,965.746	-0.34%	993,706	1,972.520	1,955.281	-0.87%	1,018,774

Notes: Comparison of actual and simulated spending, baseline structural model from Section 7 vs. richer structural model from Appendix D.2. In each comparison, data are shown only for individuals within the 1st to 99th percentiles of the distribution of the predicted trend in consumption; individuals too dissimilar from the regression sample in this dimension are dropped.

Overall, this analysis demonstrates that the simple approach, which restricts the sample to individuals on the linear portions of the budget set and ignores bunching and switching behavior, performs as well as more complex nonlinear models of consumption.

F Additional Tables

Appendix Table 14: Results of Full Year ICR (and Donut) Price Regressions, with Stark Donut Coverage and Deductible Variables – All Enrollees

Price	Full Sample, All Years Pooled			
	Including Donut		Excluding Donut	
	Coef	SE	Coef	SE
ICR	-0.078	0.006 **	-0.079	0.006 **
Donut	0.021	0.009 *		
Ded. Chg	-0.039	0.006 **	-0.040	0.006 **
Stark	-0.043	0.008 **	-0.037	0.008 **

Notes: Results of full year regressions of log consumption change on log change in ICR and donut prices, as well as changes in stark gap coverage and deductible coverage. Full sample of enrollees included. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

Appendix Table 15: Results of Full Year ICR (and Donut) Price Regressions, with Stark Donut Coverage and Deductible Variables – All Enrollees vs. Non-Switchers Only

Price	Full Sample, All Years Pooled			
	All N=2,707,315		Non-Switchers N=2,435,952	
	Coef	SE	Coef	SE
ICR	-0.078	0.006 **	-0.077	0.006 **
Donut	0.021	0.009 *	0.022	0.009 *
Ded. Chg	-0.039	0.006 **	-0.036	0.005 **
Stark	-0.043	0.008 **	-0.029	0.008 **

Notes: Results of full year regressions of log consumption change on log change in ICR and donut prices, as well as changes in stark gap coverage and deductible coverage. Full sample of enrollees included. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

Appendix Table 16: Results of Quarterly ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – High- and Low-Spending Enrollees, Pooled All Years, Pooled vs. Separate Regressions

Period	Price	Pooled Regression				Separate Regressions			
		Low-Spending Enrollee Response		High-Spending Enrollee Response		Low-Spending Enrollee Response		High-Spending Enrollee Response	
		Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	-0.054	0.006 **	-0.027	0.008 **	-0.062	0.006 **	-0.026	0.008 **
Q1	Donut	0.023	0.007 **	-0.045	0.010 **	0.035	0.008 **	-0.061	0.012 **
Q1	Ded. Chg	-0.048	0.006 **	-0.048	0.006 **	-0.033	0.005 **	-0.042	0.006 **
Q1	Stark	-0.016	0.009	-0.016	0.009	-0.025	0.007 **	-0.003	0.008
Q2	ICR	-0.046	0.006 **	-0.010	0.009	-0.054	0.006 **	-0.016	0.008
Q2	Donut	0.029	0.007 **	-0.037	0.014 **	0.034	0.007 **	-0.088	0.014 **
Q2	Ded. Chg	-0.024	0.005 **	-0.024	0.005 **	-0.014	0.006 *	-0.004	0.007
Q2	Stark	0.000	0.009	0.000	0.009	-0.003	0.008	0.006	0.011
Q3	ICR	-0.045	0.006 **	-0.003	0.009	-0.047	0.006 **	-0.015	0.009
Q3	Donut	0.015	0.008	-0.076	0.013 **	0.010	0.008	-0.109	0.015 **
Q3	Ded. Chg	-0.016	0.006 **	-0.016	0.006 **	-0.011	0.006	-0.004	0.008
Q3	Stark	-0.008	0.009	-0.008	0.009	0.006	0.008	-0.011	0.010
Q4	ICR	-0.060	0.008 **	0.022	0.012	-0.060	0.008 **	0.009	0.011
Q4	Donut	0.008	0.010	-0.162	0.017 **	0.001	0.010	-0.204	0.020 **
Q4	Ded. Chg	-0.012	0.007	-0.012	0.007	-0.005	0.006	-0.005	0.010
Q4	Stark	-0.047	0.011 **	-0.047	0.011 **	-0.046	0.010 **	-0.030	0.013 *

Appendix Table 17: Estimated Structural Model Parameters, Pooled All Years – By Chronic Condition

Description	Full Sample [N=1,985,229]		Chronic (All) [N=1,330,677]		Non-Chronic [N=627,941]	
	Structural Parameter	Mean Estimate	Structural Parameter	Mean Estimate	Structural Parameter	Mean Estimate
Days supply (P=0)	α	371.356	α	388.956	α	331.509
Myopia	β	0.321	β	0.394	β	0.303
Marginal price effect	η	-1.886	η	-1.958	η	-1.779
Deductible effect	K_{deduct}	-8.830	K_{deduct}	-9.113	K_{deduct}	-7.703
Stark gap effect	K_{stark}	-6.358	K_{stark}	-5.801	K_{stark}	-13.634
Implied elasticity	ϵ	-0.117	ϵ	-0.121	ϵ	-0.109

Description	Hypertension [N=843,867]		Hypercholesterolemia [N=338,802]		Diabetes [N=308,165]	
	Structural Parameter	Mean Estimate	Structural Parameter	Mean Estimate	Structural Parameter	Mean Estimate
Days supply (P=0)	α	390.579	α	359.010	α	476.529
Myopia	β	0.377	β	0.469	β	0.509
Marginal price effect	η	-2.017	η	-1.784	η	-2.018
Deductible effect	K_{deduct}	-10.199	K_{deduct}	-12.813	K_{deduct}	-9.554
Stark gap effect	K_{stark}	-5.168	K_{stark}	-12.454	K_{stark}	-3.086
Implied elasticity	ϵ	-0.117	ϵ	-0.128	ϵ	-0.115

Notes: Authors' calculations. Data are shown only for individuals within the 1st to 99th percentiles of the distribution of the predicted trend in consumption; individuals too dissimilar from the regression sample in this dimension are dropped. All parameters shown are averages except for the hyperbolic discount factor. Following Goldman, et al. (2004), chronic illnesses are identified using diagnosis codes from the individuals' medical claims histories.