Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets[1]

Jason Abaluck
Yale and NBER

Jonathan Gruber
MIT and NBER

Ashley Swanson
University of Pennsylvania

May 2014

[1] Preliminary and incomplete. Please do not cite or distribute without authors' permission. Kathleen Easterbrook, Ayesha Mahmud, and Adrienne Sabety provided outstanding research assistance. Financial support was provided by the National Institute of Aging.

Abstract

Medicare Part D enrollees deciding which drugs to consume face a complicated decision problem: they must choose what basket of drugs to consume in each period given potentially dozens of close substitutes with differing and often hard-to-find prices, a non-linear budget set, and a dynamic environment in which drug needs are realized gradually over time. We estimate a model of prescription drug consumption which accounts for non-linear budget sets, dynamic incentives due to myopia and uncertainty, and price salience using a 20% sample of the entire universe of Part D claims data from 2006-2009. We "linearize" the analysis by using variation away from kink points to identify underlying structural parameters; this permits estimation of a more flexible model. Dynamics in estimated price responses identify the degree to which individuals are myopic and respond to current prices versus marginal prices. We also consider whether individuals respond directly to salient plan cost-sharing characteristics after controlling for the prices they face as a result of those characteristics. Our identification strategy relies on inertia in plan choices – we estimate individuals' responses to plan cost-sharing changes between years. Our estimates suggest small marginal price elasticities and substantial myopia. We also find evidence that plan donut coverage impacts consumption more than would be expected given its impact on expected marginal price, suggesting that price salience is an important factor in determining consumption.

1 Introduction

Under the Medicare Part D prescription drug benefit, private insurers offer a wide range of products with varying prices and product features. Two key features of Medicare Part D are that plans' cost-sharing characteristics are allowed to vary along a number of dimensions so long as they are at least as generous as a standard plan set by the Centers for Medicare and Medicaid Services (CMS), and that Part D plans generally have more nonlinear pricing features than are generally observed in insurance plans. This project considers how prescription drug utilization of Medicare Part D enrollees responds to cost-sharing characteristics. As with most insurance plans, the standard plan set by CMS was designed with the goal of reducing moral hazard while providing financial risk protection for the elderly plan enrollees. Over 90% of enrollees are actually enrolled in non-standard plans, so it is critical to understand how the observed variations – elimination of deductibles, pricing tiers, coverage in the infamous Part D "donut hole," etc. – impact utilization. As a matter of policy, the donut hole will be filled in by 2020 under the Patient Protection and Affordable Care Act (ACA) of 2010; this research will inform us as to what effect this will have on spending growth, all else equal.

The complex nature of Part D plans implies that enrollees, in choosing their prescription drug utilization, face a complicated decision problem. They must choose what basket of drugs to consume in each period given potentially dozens of close substitutes, whose prices are difficult to determine without close examination of plan benefits and formularies, a non-linear budget set in which out-of-pocket drug costs for current and future periods change as a function of total spending to date, and 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 project, we present a model of prescription drug consumption which accounts for optimization with non-linear budget sets, dynamic incentives due to myopia and uncertainty, and variation in salience of different types of prices.

The main data source for this research is a 20% sample of Medicare Part D claims provided by CMS. This claims data includes information on drugs consumed (at the national drug code (NDC) level), as well as the date, quantity, wholesale price and amount paid by insurer and beneficiary for each claim. Our identification strategy utilizes the substantial inertia in plan choice present in the Medicare beneficiary population – once an initial plan has been chosen, the vast majority of enrollees remain in that plan in subsequent years even though plan cost-sharing characteristics may change substantially. We therefore analyze how year-to-year changes in cost-sharing features of plans impact changes in the pattern of prescription drug utilization – details on implementation are in Section 3 below.

We begin our analysis with a differences-in-differences regression to illustrate the thought experiment behind our identification strategy. We regress consumption on gap coverage among individuals in plans matched on all coverage features except for the donut hole. These reduced form results demonstrate that Part D enrollees' consumption is sensitive to the presence of gap coverage, all else equal, and that that sensitivity increases throughout the year as more individuals enter the donut hole.

Using a simple model of dynamic consumption with nonlinear budget sets, we then demonstrate that linear regression methods can be used to recover structural parameters for individuals whose marginal prices are in the interior of budget set segments, and these parameters can in turn be used to forecast behavior at kink points.[1] We demonstrate 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. Using a simple reduced form model in which consumers respond both to current and expected marginal prices, the estimated coverage phase price responses reveal consumers' price elasticities and degree of myopia. We find overall price elasticities of around -0.1 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 0.19, suggesting a very high

[1] Relative to the structural estimation methods commonly used in the nonlinear budget set literature (Kowalski (2012) and Einav, Finkelstein, and Schrimpf (2013) are two excellent recent examples), this approach is computationally simple and requires few functional form assumptions to recover price elasticities.

degree of myopia, again consistent with the existing literature.

We also examine whether individuals respond excessively to particularly salient plan cost-sharing features. We find evidence 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.

Using the structural parameters implied by our linear estimates, we show that it matters not just what prices are, but also when they are encountered in the year and how they are presented to enrollees. 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.

The rest of the paper proceeds as follows. In Section 2, we describe the background of the Part D program and the literature on decision-making among the elderly and moral hazard in health care. Section 3 describes our identification strategy. Section 4 describes the data and provides details on price variable construction. Section 5 presents our reduced form analysis of the impact of donut hole coverage on consumption. In Section 6, we lay out and estimate a structural model which estimates how consumption responds to prices allowing for both myopia and salience in a dynamic setting with uncertainty. Section 7 translates our price coefficients into structural parameters. Section 8 concludes.

2 Background on Medicare Part D, Elder Decision-Making, and Moral Hazard in the Medical Context

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.[2] The most noticeable innovation of the Part D plan is that this new Medicare benefit is not delivered by the government,

[2] Duggan 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.

but rather by private insurers under contract with the government. Beneficiaries can choose from three types of private insurance plans coverage of their drug expenditures. The first type are stand-alone plans called Medicare Prescription Drug Plans (PDP) (plans that just offers prescription drug benefits). For example, in 2006, there were 1,429 total PDPs offered throughout the nation, with most states offering about forty PDPs. The second alternative is Medicare Advantage (MA) plans, plans that provide all Medicare benefits, including prescription drugs, such as HMO, PPO, or private FFS plans. There were 1,314 total 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.

Under Part D, recipients are entitled to basic coverage of prescription drugs by a plan with a structure 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 \$2,250 of drug spending (up to \$2,500 total); 0% of costs for the next \$3,600 of drug spending (up to \$5,100 total); 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 ensure. 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 coverage gap. The ACA mandates that the donut hole be "filled in" gradually by 2020. For the 2014 benefit year, enrollees in plans that have a coverage gap are entitled to a 52.5% discount on branded drugs and a 21% discount on generics while in the coverage gap.[3]

Enrollment in Part D plans was voluntary for Medicare eligible citizens. In order to

[3]Drug manufacturers offer the branded discount under the Medicare Coverage Gap Discount Program; Medicare covers the 21% generic discount in the gap (CMS, 2013).

mitigate adverse selection, Medicare recipients not signing during the initial enrollment period in the first year of the program or when they aged into Medicare (and who did not have other creditable prescription drug coverage) were subject to a financial penalty if they eventually joined the program.[4]

Our project builds on several literatures. First, we consider decision-making in a complex setting by an elderly population. Issues considered in behavioral economics, such as myopia and salience, may be paramount within the context of the elderly, given that the potential for cognitive failures rises at older ages. A recent study by Sumit Agarwal, John C. Driscoll, Xavier Gabaix, and David Laibson (2006) 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.[5] Several studies of these issues apply specifically to the context of Part D. For example, Florian Heiss, Daniel McFadden and Joachim Winter (2006); Jeffrey R. Kling, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen and Marian V. Wrobel (2008); Abaluck and Gruber (2011); and Ketcham et al. (2011) each study plan choice under Medicare Part D and find striking evidence in a variety of empirical settings that elders do not make cost-minimizing choices of Part D plans (though there is some disagreement regarding whether choices improved over time). Our project 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.

There is also a rich literature on the impact of cost-sharing on health care utilization utilization and this literature is reviewed in great detail in Chandra, Gruber, and McKnight (2008). Of particular note is the RAND Health Insurance Experiment (HIE), which

[4]One 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 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.

[5]See also Timothy A. Salthouse (2004), which shows clear evidence that the performance on a series of memory and analytic tasks declines sharply after age 60; and Laura Fratiglioni, Diane De Ronchi, and Hedda A. Torres (1999) for evidence on the relationship between the onset of dementia and age.

is summarized in Manning et al. (1987) and Newhouse (1993). 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 (2007) estimate an arc-elasticity of -0.47. Chandra, Gruber, and McKnight (2008) 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 (2011) studies the aggregate utilization of 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 deadweight losses from moral hazard. Aron-Dine, Einav, Finkelstein, and Cullen (2012) 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 (2013) 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. They estimate a weekly β of 0.93, which translates roughly to a quarterly β of 0.42; they find static price elasticities ranging from -0.45 to -0.75. Our strategy builds on this literature to estimate elasticities with respect to variation in both current and marginal price for a broad range of the overall spending and age distributions.

3 Identification Strategy

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 price (OOP) and consumption, since sicker individuals are more likely to end up in the donut hole, all else equal.

To deal with the first and second issues, we instrument for prices using variation generated by changes in Part D plan cost sharing rules and by taking advantage of the substantial inertia in Part D plan enrollment. The thought experiment that motivates this strategy is as follows. Consider two elderly individuals, Sheldon and Leonard, who choose their plans in 2006 and plan to stay in that plan for several years before they reoptimize. They choose different plans in 2006, but these plans have identical cost-sharing provisions, and Sheldon and Leonard use identical prescription drugs in 2006. In 2007, Plan A, in which

Sheldon is enrolled, increases its copayments, while Plan B, in which Leonard is enrolled, does not. Since neither Sheldon nor Leonard is reoptimizing, there is an exogenous shift in cost sharing between them from 2006 to 2007. That is, any difference in spending in 2007 between Sheldon and Leonard is due to cost sharing changes rather than active plan choices.

Of course, to the extent that we don't see Sheldon or Leonard switch plans, we can't say for certain whether this is because of inertia or because of preferences; it may be that Sheldon stayed in plan A not because of a failure to reoptimize, but precisely because he anticipated having lower prescription drug needs next year, which would lead to the same endogeneity bias noted above. However, we can compare all individuals who were in plan A in 2006 to all individuals who were in plan B in 2006, conditioning on any differences in characteristics between these two groups (including differences in 2006 utilization). So long as there is some inertia in plan choice, then, on average, individuals who were in plan A in 2006 will see higher copayments in 2007 than those who were in plan B in 2006. That is, if some individuals don't reoptimize for 2007, there is an exogenous change on average in copayments for the entire group enrolled in plan A in 2006. Given the small degree of switching observed in practice (about 10% in each pair of years in our sample), it seems likely that many individuals are not annually reoptimizing, a conclusion strongly supported by Abaluck and Gruber (2013).

The standard approaches to the third problem in the literature are to either estimate a nonlinear structural model assuming a particular model of optimization behavior as in Hausman (1985) or, more recently, Kowalski (2011) and Einav, Finkelstein, and Schrimpf (2013); or to estimate a nonparametric model with higher order terms for each segment and threshold of the budget set as in Newey and Blomquist (1995). In our analyses, we employ a simplified version of Newey and Blomquist by considering the linear response to budget set segments (there is no meaningful cross-sectional variation in threshold locations) 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 show similar patterns.

Our identification approach contains 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 that this allows us to identify a marginal price response for nearly the full scope of the enrollee population 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. (2012) 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.

4 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 (days supply – this is our outcome variable in all specifications) consumed, 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.

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) or 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; sample period 2006-7 included 1,372 plans, while 2007-8 and 2008-9 each included over 1,700 plans. Summary statistics on sample plans and enrollees are shown in Table 1. The majority of sample enrollees are white and female, with a mean age of 75. Between the first and second year of each year pair, a small proportion (9-11%) of enrollees switched plans.

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. 18-24% of enrollees were in plans with the standard deductible, but 70-80% of enrollees were in plans with no deductible, and a small fringe of enrollees were in plans with positive, but nonstandard, deductibles. Furthermore, a nontrivial proportion of enrollees had coverage in the donut hole; in the early years of the program, 1-6% of enrollees had full donut coverage, and 13-20% of enrollees had some coverage in the donut hole throughout our sample period.

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 sample enrollee in later year pairs, as would be 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. For year pair $(2006, 2007)$, 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

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%
% No Donut	6.09%	2.18%	1.43%	0.00%	0.00%	0.00%
% Gap 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.

enrolled in plan p in coverage phase c, year y of year pair $year1$-$year2$, is defined as

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

where $CS_{dcy,p}$ and $CP_{dcy,p}$ are coinsurance rates and copays, respectively, for drug d in plan p, $RP_{dcy,p}$ is retail price for drug d in plan p, and the consumption weight for drug d is calculated as

$$w_{idy} = q_{id,y} / \sum_{d \in D_i} q_{id,y}$$

using observed quantity consumed $q_{id,y}$ for each individual-drug-year combination. D_i is the set of all drugs consumed by individual i.

The price *instrument* in phase c of year y is the weighted average price the individual would face in phase c given the 2006 plan's year y cost-sharing rules; weights use the individual's 2006 consumption. For coinsurances, we apply the coinsurance rate to the retail price appropriate for the given plan-drug combination. The given plan-drug's average retail price in year y observed in the claims is applied for actual prices; the plan-drug's average retail price in 2006 is applied for price instruments. Prices are for 30-day supplies of drugs; when 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. That is, the IV price P^{IV}_{icy} is defined as

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

Variation in instrumental variables (IVs) in first year of each year pair (left panel) and differences across years (right panel) are shown in Table 2. 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.

Note that, for brevity, price differences are shown only for the ICR and donut hole;

individuals with deductibles in both years of the year pair have no effective IV variation 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.[6]

In looking across the three year pairs of our analysis sample, we note several patterns of interest. First, there is substantial variation in the year 1 donut price IV in each year pair, but that variation is driven primarily by variation across individuals in the retail prices of the drugs they take, as the majority of the sample has no gap coverage. Our analysis is done on prices in differences, and we see that the variation in the IV donut price *difference* is falling over time. The IV prices defined above use year 1 retail prices and consumption weights, so the donut IV is only nonzero for enrollees with some gap coverage in at least one year; as we saw in Table 1, fewer enrollees have gap coverage over time. Second, the IV and donut price differences have both negative and positive values, but are skewed positive, so that the average enrollee experiences diminishing plan generosity between years 1 and 2.

5 Illustrative Example – Gap Coverage

As an illustrative example of our identification strategy, we start by considering how individuals respond to gap coverage changes, all else equal. To isolate the consumption response to the gap 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 gap coverage in 2006, and we impose that matches have similar prices in year 1 and pre-gap in year 2 and different prices in the gap in year 2.[7] Taking unique plan matches only, we obtain

[6] In some analyses below, we also analyze responses to variation coming from changes in the location of deductible thresholds between years.

[7] In 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

Table 2: Sample Price Instrument Variation

	Price Year 1				Price Year 2 - Price Year 1	
	Ded	ICR	Donut	Catas	ICR	Donut
	Year Pair: 2006-2007					
N	451,632	451,632	451,632	451,632	451,632	95,537
5th Prctile	5.55	4.83	12.65	2.07	-4.12	-15.84
Median	20.16	15.23	45.46	3.80	0.00	0.00
95th Prctile	69.82	31.56	98.46	5.73	8.51	46.71
Mean	26.81	16.58	50.80	3.99	1.01	5.64
Std. Dev.	31.43	14.73	60.12	2.98	7.93	36.77
	Year Pair: 2007-2008					
N	1,126,682	1,126,682	1,126,682	1,126,682	1,126,682	187,139
5th Prctile	5.00	3.68	10.00	2.15	-3.86	-5.00
Median	16.83	13.62	42.60	3.69	0.58	0.00
95th Prctile	63.29	30.80	101.99	5.93	9.00	32.37
Mean	23.04	15.41	49.62	3.95	1.32	3.49
Std. Dev.	34.65	15.76	60.02	2.89	7.10	20.04
	Year Pair: 2008-2009					
N	1,129,200	1,129,200	1,129,200	1,129,200	1,129,200	159,486
5th Prctile	4.00	2.79	9.23	2.25	-1.98	-1.04
Median	14.79	12.02	40.38	3.52	1.89	0.00
95th Prctile	62.95	31.52	106.84	6.14	12.75	6.07
Mean	21.88	14.50	49.19	3.90	3.19	1.28
Std. Dev.	44.22	28.47	98.43	4.83	7.33	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 (price instrument difference is by definition zero for other plans).

a large sample of approximately 100,000 individuals in matched plans, where plans that either dropped generic coverage between 2006 and 2007 or added it were matched to plans

restrictive matches have larger standard errors.

Figure 3: Matched Plans – OOP Cost Comparison (2007)

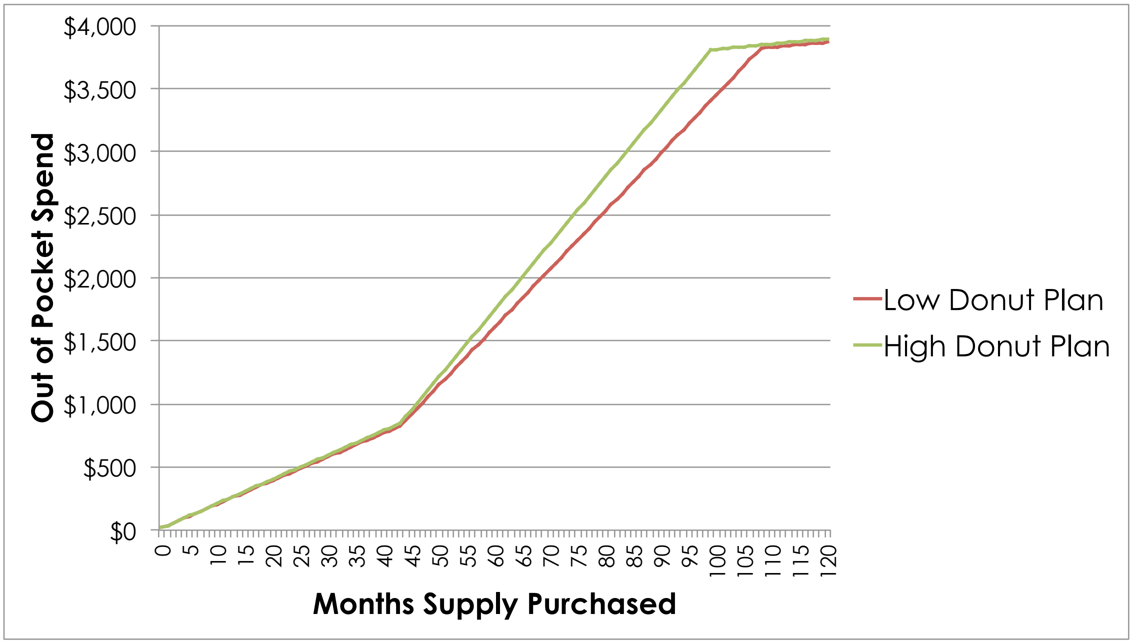

which did not change their gap coverage between years.[8] See Figure 3 for a comparison of out of pocket cost as a function of days supply purchased for "high donut" vs. "low donut" matched plans. None of the matched plans had deductibles, and we see that total out of pocket cost is equivalent in the 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 rapidly.[9]

Using this sample, we run 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 in the "high gap price" plan in the second year (2007). These regressions include con-

[8]Many 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 2007 prices and least similar 2007 donut prices) for the largest matched plan, 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.

[9]The 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.

trols for a rich set of individual demographics and individual drug 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, including 50 quantiles each of days supply purchased, total drug expenditure, and out-of-pocket (OOP) drug expenditure. We also include dummies for having any drug spending in each generic therapeutic class (GTC), which is a summary measure of target medical condition.[10] 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. 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).

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

We see several striking facts in Table 4. First, enrollees' consumption responds negatively to lower gap coverage. Second, the response is much stronger at the end of the year, when people are more likely to have entered the donut hole. Myopia is one possible explanation for this pattern – in the next section we present more direct evidence that it plays a substantial role.[11]

[10] Each 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.

[11] Note that we do not necessarily expect the full year coefficients to be a mechanical average of the quarterly coefficients, 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.

Table 4: Differences-in-Differences Results, Gap Coverage Dummy, Matched Plan Enrollees

Price	Obs	Coef	SE
Full Year	97,330	-0.020	0.013
Q1	78,867	0.001	0.014
Q2	78,867	-0.025	0.012 *
Q3	78,867	-0.039	0.012 **
Q4	78,867	-0.068	0.013 **

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.

6 Model and Results

We begin with a simple model, in which forward-looking, rational individuals optimize their consumption of prescription drugs over the course of the year with no uncertainty. Suppose individuals maximize utility $u(q)$ by choosing consumption q_t for each $t = 1, ..., T$ and suppose that they face a nonlinear budget set with the following expenditure function over total quantity purchased in the current year:

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

with out-of-pocket prices $p_1 < p_2$ as would occur at the convex Medicare Part D donut threshold $\bar{X}$ and retail price (total price paid by plan *plus* enrollee) R. In 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 $\bar{X} = \$2,500$.

Given perfectly forward-looking behavior and no uncertainty, the dynamic problem collapses to the static choice of prescription drug quantity for the full year, and for a general quasiconcave, continuously differentiable utility function over prescription drugs $u(q)$ (and assuming linear utility for the numeraire) we find

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

so that individuals consume either at the linear optimum for p_1 or p_2 or "bunch" at the coverage threshold $\bar{X}$. This 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 reduce their consumption even more under the higher linear price p_2, so their consumption is 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). Thus, for many price-parameter combinations, this model predicts that total consumption will be a function of marginal (end of year) price as in a linear demand model. Those consumers who would be willing to pay p_1 for an additional unit of consumption at the kink but are not willing to pay the post-kink price p_2 for that unit will bunch at the donut hole.[12]

We begin by empirically estimating a special case of this model with power utility (motivating a log functional form) and ignoring coverage range switching (so this model applies within coverage ranges but will not be correct insofar as changes in prices induce individuals to move between coverage ranges or bunch at kink points). This yields an equation of the form: $log(q_{it}) = \alpha_0 + \alpha_1 log(MP_{i,y})$ where $MP_{i,y}$ is the marginal end of year price for individual i in year y. For each pair of years, we take differences across years,

[12] 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.

yielding an equation of the form:

$$log(q_{i,y}) - log(q_{i,y-1}) = \alpha_0 + \alpha_1(log(MP_{i,y}) - log(MP_{i,y-1})) + u_{iy}$$

As described in Section 4, 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.

Expected end-of-year price is constructed by taking, for each sample individual i, a set of 200 other sample individuals in the same centile of spending, running those individuals' claims through individual i's plan characteristics, determining the marginal coverage phase given those individuals' claims, and taking a weighted average over the marginal coverage phases for those 200 individuals to determine i's expected end of year price.[13] 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 – 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 control for dummies for year 1 deductible and gap 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.[14] The results of these regressions for 2006-2007, 2007-2008, and 2008-2009 are displayed in Table 5 below:

[13] For actual prices, centile of spending is based on the contemporaneous year's spending; for the year 2 price IV, year 1 spending centile is used.

[14] In the paired analysis in the previous section, the plan characteristics would be collinear with the plan pair fixed effects and "high donut" dummy. In this analysis, we can alternatively exclude plan characteristics and include plan fixed effects – the results are unchanged by this alternative specification.

Table 5: Results of Full Year Marginal Price Regressions, All Enrollees

	OLS		Reduced Form		First Stage		IV	
Years	Coef	SE	Coef	SE	Coef	SE	Coef	SE
2006-7	0.049	0.001 **	-0.047	0.005 **	0.558	0.006 **	-0.085	0.009
2007-8	0.031	0.001 **	-0.027	0.003 **	0.419	0.004 **	-0.064	0.008
2008-9	0.015	0.001 **	-0.037	0.003 **	0.604	0.004 **	-0.062	0.005

Notes: Results of full year marginal price regressions with full controls. Superscript ** indicates significance at the 1% level; superscript * indicates significance at the 5% level.

We first note that the OLS specification is wrong-signed in the sense that higher prices are associated with higher consumption. This is consistent with the several problematic forms of endogeneity mentioned above; for example, suppose that sicker individuals have steeper increasing trends in drug consumption. If the only price change experienced by Part D enrollees were based on drug retail prices increasing year to year, then we would find a mechanical positive relationship between consumption change and price change, as sicker individuals who consume into the donut hole have higher marginal price changes and higher consumption changes than healthier individuals who end the year in the ICR. We also note that our first stage is very strong - IV price changes are highly predictive of actual price changes based on cost-sharing changes alone, though our identification strategy of shutting down retail price and consumption-based price changes leads the IV price change to increase less than one-for-one with the actual price change.

Our reduced form and IV estimations are right-signed and imply a small but significant elasticity of drug consumption. We find that total consumption is responsive to expected marginal price in all year pairs, and find similar magnitudes across years, suggesting an elasticity of prescription drug consumption of about -0.06 to -0.08. This estimate is smaller than the medical care utilization elasticity of -0.2 found in the RAND HIE, but similar to the prescription drug utilization elasticities found among PPO enrollees by Chandra, Gruber, and McKnight (2010).

These results suggest that this simple regression specification provides a reasonable picture of prescription drug utilization in the presence of a nonlinear budget set. However, they ignore several factors which may be important. First, they assume perfectly forward-looking behavior and no uncertainty. Second, they ignore the nonlinearity of the budget set and accordingly will be biased to the extent that a non-negligible proportion of the sample "bunches" at the donut coverage threshold or are induced to switch marginal coverage phase by the price change between years 1 and 2.[15] Third, they assume perfect rationality and perfect observability of prices, which, in the case of elder decision-making with respect to a complicated price schedule, may be overly restrictive.[16] While remaining in this reduced-form framework, we use several additional analyses to allow for dynamic decision-making, distinguish myopia from decision-making under uncertainty, and allow for some prices to be more salient than others.

6.1 Dynamic Decision-Making with Myopia

First, we relax the above model to allow for dynamic decision-making with myopia. Assume that individuals make their prescription drug utilization choices in each of T periods, and that utility is given by:

$$\begin{aligned} V_T(X) &= W_T(X) = max_q u(q) - p(X,q) \\ V_t(X) &= max_q u(q) - p(X,q) + \beta W_{t+1}(X + p(X,q)) \\ W_t(X) &= max_c u(q) - p(X,q) + W_{t+t}(X + p(X,q)) \end{aligned}$$

where $u(q)$ is quadratic. That is, we allow for hyperbolic discounting as in a $\beta - \delta$ discounting model, but impose that $\delta = 1$.[17] This yields a specification in which, for all individuals

[15] For example, suppose that a given plan increases the price of drugs in the ICR but decreases the price of drugs in the donut hole and in doing so induces a sample of individuals who were formerly below the donut threshold to enter into the donut. Our identification strategy would suggest consumers in the ICR increased their consumption in response to the price increase in the ICR (a wrong positive elasticity) when in fact they increased consumption in response to the price decline in the donut hole.

[16] See, e.g., Abaluck and Gruber (2011), which considers the case of inconsistencies in Medicare Part D prescription drug plan choice.

[17] Given that all dynamic decisions are made within the relatively short time horizon of a single year (Part D contracts exist for a maximum of one year, and the only dynamic consideration in this simple model is the effect of current consumption on future price within the current year's contract), it seems reasonable to

ending both the current period and the year in the interior of a coverage phase, consumption in each period is given by: $q_{it} = \alpha_0 + \alpha_1(\beta MP_{i,y} + (1-\beta)CP_{it,y})$ where $MP_{i,y}$ is the marginal end of year price and $CP_{it,y}$ is the current price. This parallels the current-future price specification in Aron-Dine et. al. (2012), although as noted above, we can separately estimate both current and future price elasticities.

In practice, given that current and future prices may be complicated objects for individuals to calculate, particularly in the presence of behavioral biases, we consider a more flexible specification of the dynamic utilization model which should subsume the current/future model without imposing a specific functional form on how individuals respond to different segments of the budget set. In particular, we regress quarterly quantity consumed on the initial coverage range (ICR) price and the donut hole price.[18] In the rational model in which consumers respond only to marginal end of year prices and know exactly what these prices will be, this model predicts that the respective coefficients should scale with the probability that each coverage range is marginal. That is, we can write:

$$\begin{aligned} q_{it} &= \alpha_0 + \alpha_1 MP_{i,y} \\ &= \alpha_0 + \alpha_1 \mathbb{1}(m_i = ICR)P_{ICR} + \alpha_1 \mathbb{1}(m_i = Donut)P_{Donut} \end{aligned}$$

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

$$q_{it} = \alpha_0 + \alpha_{ICR}P_{ICR} + \alpha_{Donut}P_{Donut},$$

the estimated coefficients should then be such that $E(\alpha_{ICR}) = E(\alpha_1 \mathbb{1}(m_i = ICR)) =$

assume an exponential discount factor of 1. We 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. For example, as in the "spotlighting" version of "schmeduling" consumption as discussed in the context of nonlinear price schedules in Liebman and Zeckhauser (2004) – consumers may respond to immediate or local prices and ignore the full schedule that they face.

[18] The 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 order to control for potential correlation between the initial coverage range and donut prices and deductible coverage, we also include controls for the year 2 deductible threshold of the year 1 chosen plan.

$\alpha_1 \Pr(m_i = ICR)$ and $E(\alpha_{Donut}) = E(\alpha_1 \mathbb{1}(m_i = Donut)) = \alpha_1 \Pr(m_i = Donut)$.[19] In the linear demand specification (quadratic utility), the same result holds even if there is uncertainty about which coverage range is marginal as long as price changes do not alter the respective probabilities of ending in each coverage range (the analogue of our assumption above that price changes do not induce switching). For 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.[20]

In the more general current/future specification considered above, we have:

$$\begin{aligned}
q_{it} &= \alpha_0 + \alpha_1(\beta MP_{it,y} + (1-\beta)CP_{i,y}) \\
&= \alpha_0 + \alpha_1\beta\mathbb{1}(m_i = ICR)P_{ICR} + \alpha_1\beta\mathbb{1}(m_i = Donut)P_{Donut} \\
&+ \alpha_1(1-\beta)\mathbb{1}(c_i = ICR)P_{ICR} + \alpha_1(1-\beta)\mathbb{1}(c_i = Donut)P_{Donut} \\
&= \alpha_0 + \alpha_{ICR}P_{ICR} + \alpha_{Donut}P_{Donut},
\end{aligned}$$

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 should be $E(\alpha_{ICR}) = \alpha_1(\beta \Pr(m_i = ICR) + (1-\beta)\Pr(c_i = ICR))$ and $E(\alpha_{donut}) = \alpha_1(\beta P(m_i = Donut) + (1-\beta)P(c_i = Donut)$. In other words, the results of this 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 consumption response be scaled as a

[19] We are implicitly assuming here that the LATE estimated given our instrument will equal the ATE. This assumption is equivalent to asking whether compilers - 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.

[20] In the log model this result no longer holds exactly. $log(C_i) = \alpha_0 + \alpha_1 log(E_i(P_i))$ 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. Assuming as above that price changes do not alter the respective probabilities of ending in each coverage range, the elasticity with respect to each coverage range price is given by: $\frac{\partial logC}{\partial log P_c} = \frac{\alpha_1 P(\text{coverage range}=c)P_c}{\sum_c P(\text{coverage range}=c)P_c}$ for everyone. Note that if all coverage ranges have the same prices, this reduces to $\alpha_1 P(\text{coverage range} = c)$, which is exactly the result we get with no uncertainty. In reality, prices vary across coverage ranges, so in future results we will include coefficients adjusted by the scaling factor $\frac{P_c}{\sum_c P(\text{coverage range}=c)P_c}$ and show that adjusting for it makes little difference for our conclusions.

function of current and future weighting under a specific model of expectations.

6.2 Quarterly Results for Individuals in Interior of Coverage Phases

The dynamic specification outlined above only holds for individuals ending both the current and marginal period in the interior of a coverage phase. In this section, we describe our first empirical specification, which models dynamic (quarterly) consumption as a function of ICR and donut prices for just such individuals – individuals for whom the donut is marginal and individuals for whom the ICR is marginal (excluding individuals likely to be on the margin between the two groups). Specifically, we look only at individuals consuming less than or equal to $1,500 in the first year of each year pair and, separately, at individuals consuming between $3,000 and $5,000 in the first year of each year pair. The first group is a set of individuals who are almost certainly going to have the ICR price be their marginal price; the second group is a set of individuals who are almost certainly going to have the donut price as their marginal price.[21]

Consider first the low-spending group. Results are in Table 6. 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%) as the small proportion of individuals with deductibles enter the ICR. 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 pooled results show that the ICR response is slightly increasing as more individuals enter their marginal coverage phase; among these low-spending enrollees, this pattern is consistent with 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.[22] The magnitude of the static price elasticity is on the lower end

[21] Of course, to the extent that prescription drug needs evolve over time, this approach is imperfect; for example, 7% of individuals in the low-spending group in 2006 reach the donut hole in 2007 and 12% of individuals in the high-spending group do not. However, we can use these samples to examine behavior in the presence of far more limited marginal price uncertainty and scope for switching than in the previous results.

[22] There is a small positive response to the donut price in the pooled results at the beginning of the year, which equals the largest of the responses across other years rather than being an approximate average. We were unable to include year-specific controls in the pooled regressions prior to this draft; the separate results

of the elasticities suggested by the literature (-0.05 to -0.06).

Consider next the high-spending group. Results are in Table 7. 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%, 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 such that high-spending individuals have a large and significant donut price response in quarter 4 (-0.22) which is more than three times the donut response at the beginning of the year (-0.056). This 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. We do not detect a response among enrollees to the "spot" price they face prior to the donut hole (the ICR price), whereas the current-future price model derived from hyperbolic discounting predicts that we should observe such a response.

Table 6: Results of Quarterly ICR and Donut Price Regressions – Low-Spending Group

			All Years Pooled		2006-7		2007-8		2008-9	
Period	Price	% in ICR	Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	83.0%	-0.043	0.006 **	-0.047	0.014 **	-0.022	0.010 *	-0.061	0.009 **
Q1	Donut		0.022	0.008 **	-0.008	0.023	0.022	0.011 *	0.018	0.010
Q2	ICR	92.1%	-0.045	0.006 **	-0.046	0.018 *	-0.023	0.010 *	-0.056	0.009 **
Q2	Donut		0.025	0.008 **	0.041	0.026	0.017	0.011	0.029	0.011 *
Q3	ICR	95.4%	-0.045	0.006 **	-0.057	0.017 **	-0.032	0.011 **	-0.044	0.008 **
Q3	Donut		0.003	0.009	0.026	0.027	0.000	0.013	0.003	0.012
Q4	ICR	97.0%	-0.056	0.008 **	-0.039	0.016 *	-0.045	0.015 **	-0.067	0.011 **
Q4	Donut		-0.020	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.

by year suggest that that will alleviate the problem in future results.

Table 7: Results of Quarterly ICR and Donut Price Regressions – High-Spending Group

			All Years Pooled		2006-7		2007-8		2008-9	
Period	Price	% in ICR	Coef	SE	Coef	SE	Coef	SE	Coef	SE
Q1	ICR	0.5%	-0.030	0.009 **	-0.018	0.018	-0.012	0.015	-0.057	0.017 **
Q1	Donut		-0.056	0.013 **	-0.054	0.021 **	-0.050	0.024 *	-0.088	0.030 **
Q2	ICR	22.2%	-0.025	0.009 **	-0.003	0.021	-0.013	0.015	-0.053	0.017 **
Q2	Donut		-0.084	0.013 **	-0.077	0.025 **	-0.069	0.022 **	-0.085	0.030 **
Q3	ICR	88.0%	-0.022	0.010 *	-0.005	0.022	-0.020	0.017	-0.029	0.018
Q3	Donut		-0.117	0.017 **	-0.111	0.029 **	-0.116	0.031 **	-0.069	0.028 *
Q4	ICR	94.9%	-0.001	0.012	-0.021	0.025	-0.017	0.021	0.017	0.022
Q4	Donut		-0.216	0.021 **	-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.

The above results provided estimates of a consistent and significant price elasticity of demand for prescription drugs throughout the spending distribution. We also observe strong evidence of myopic utilization behavior, so that enrollees' marginal price response is evidenced much more strongly at the end of the year, once they have entered their marginal coverage phase. In the next section, we examine whether individuals respond to "prices" other than the budget set segments that should be most relevant for them given their observed drug consumption patterns.

6.3 Salience Results

As discussed briefly in Section 1, 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.

As noted by a number of researchers, including Abaluck and Gruber (2011) and Chetty, Looney, and Kroft (2009), consumers may underreact to less salient prices and overreact to more salient prices. 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 more visible measures of plan coverage generosity, the current-future price specification that results from our hyperbolic discounting model may fail to adequately account for the full scope of individuals' price responses.

Some prices that may specifically be more salient are the presence or absence of a deductible, and the presence or lack of donut hole coverage. On the latter point, we note that, although some of the donut hole price variation experienced by our sample enrollees is of the sort of stark price variation used to generate arc elasticities of demand in the RAND experiment, in the Chandra, Gruber, and McKnight article, and in numerous other studies, some of the other donut pricing variation would be much less simple for enrollees to translate into prices. In 2006-7, all donut gap 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 gap coverage; see Table 8. For example, a large number of plans with no gap coverage in 2006 added generic gap coverage in 2007, implying large price decreases for the donut hole of about $10 per 30-day supply on average. 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 more in 2008-9 had slight alterations in their generic coverage in the donut hole which did not entail large average price changes and which were not generally easily understandable; see Table 9 and Table 10. For example, a large number of plans 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 price 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. The fact that essentially none of the donut price variation in 2008-9 is of this "stark" form (in contrast to 2006-7 and 2007-8) may account for the observation in Table 7 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.

Table 8: Enrollment by Gap Coverage Changes, 2006-7

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

Table 9: Enrollment by Gap Coverage Changes, 2007-8

2007 Gap Coverage, 2007-8 Sample	2008 Gap Coverage, 2007-8 Sample: Some Generics	All Generics and Some Brands	All Preferred Generics	All Generics	None
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

In order to account for the potential impact of more salient price changes on consumption behavior, we next perform additional regressions of change in quantity consumed on the initial coverage range price, the donut hole price, and two variables capturing changes in deductible or gap coverage. 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

Table 10: Enrollment by Gap Coverage Changes, 2008-9

2008 Gap Coverage, 2008-9 Sample	2009 Gap Coverage, 2008-9 Sample: Some Generics	All Generics	All Generics and Few Brands	Many Generics	Many Generics and Few Brands	None
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

2006 and 2007).[23] The stark gap coverage change variable 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; and 0 otherwise. The results for the full year are shown in Table 11.

Table 11 shows that Part D enrollees do respond on average to stark changes in gap coverage beyond such changes' effects on expected donut hole prices. The elasticities vary across year pairs but are negative across all years where such changes were observed for a non-negligible sample of enrollees (i.e., not 2008-9). We also observe in 2007-8 and 2008-9 a substantial effect of changing deductible coverage on consumption. The deductible is not marginal for the vast majority of enrollees, but all enrollees with deductibles do 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 than based on the implied out-of-pocket price change).

Table 12 shows the results of these same regressions separately for high- and low-

[23]For 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 11: Results of Full Year ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – All Enrollees

	All Years Pooled		2006-7		2007-8		2008-9	
Price	Coef	SE	Coef	SE	Coef	SE	Coef	SE
ICR	-0.076	0.006 **	-0.084	0.012 **	-0.066	0.011 **	-0.082	0.008 **
Donut	0.011	0.010	0.008	0.022	0.006	0.014	0.035	0.012 **
Ded. Chg	-0.039	0.007 **	-0.008	0.012	-0.039	0.009 **	-0.048	0.008 **
Stark	-0.033	0.007 **	-0.038	0.011 **	-0.047	0.017 **	0.009	0.027

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.

Table 12: Results of Full Year ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – High- and Low-Spending Enrollees, Pooled Years

	Low-Spend Enrollees		High-Spend Enrollees	
Price	Coef	SE	Coef	SE
ICR	-0.095	0.010 **	-0.022	0.007 **
Donut	0.044	0.015 **	-0.066	0.012 **
Ded. Chg	-0.048	0.010 **	-0.017	0.007 *
Stark	-0.044	0.015 **	-0.018	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. All year pairs pooled. Superscript (**) indicates significance at the 1% level; superscript (*) indicates significance at the 5% level.

spending enrollees. Interestingly, we observe that the stark gap coverage response is negative and significant *even among low-spending individuals* who have essentially zero probability of reaching the donut hole. The "stark" coverage change elasticity is smaller in magnitude for high-spending individuals; this is consistent with salience effects being less strong for individuals who have experienced donut hole prices in previous years, but the coefficients for high and low spenders are not directly comparable because consumption is in logs.

We also observe in Table 12 that the deductible change response is significantly larger for low-spending enrollees than for 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, so this is consistent with myopic response to the deductible variation; however, as noted above, high- and low-spending enrollees price responses are not directly comparable.[24]

Taken together, these results imply 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, there is also a significant response to stark changes in gap coverage 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.[25]

7 Structural Interpretation

Our results thus far have provided evidence of significant price responses overall and substantial myopia; we now relate those regression estimates to the structural parameters from the dynamic consumption model in Section 6.1. This permits us to construct counterfactual estimates for beneficiaries nearer the kink point. Because our estimates are from a log-log regression specification and our model of hyperbolic discounting assumes a quadratic util-

[24]See Table 13 for further evidence that deductible response is evidence of myopia – deductible responses are stronger at the beginning of the year than at the end of the year.

[25]The ICR-donut regression results alone may not capture some rational responses to coverage changes in the presence of uncertainty, because our price variables pertain for each individual based on their predicted spending. For example, if low-spending individuals consume generic drugs only and are enrolled in plans that move from "All Branded and Generic" coverage in the donut hole to "Generic" coverage only, their donut hole price change in the regression results above would equal zero and we would not observe any true consumption change with respect to this substantial change in plan generosity. However, the low uncertainty in marginal coverage phase for our inframarginal individuals suggests that these results are evidence of price salience effects given any model of rational expectations.

ity function (linear demand), this exercise requires that we linearize our main empirical specification (including the deductible and stark gap coverage terms):

$$\begin{aligned}
log(q_{i,y}) - log(q_{i,y-1}) &= a + X_i * \delta + b_{ICR} * (log(P_{ICR,y}) - log(P_{ICR,y-1})) \\
&+ b_{Donut} * (log(P_{Donut,y}) - log(P_{Donut,y-1})) \\
&+ \theta_{Ded} * dedchg_i + \theta_{stark} * stark_i + u_{iy}
\end{aligned}$$

The structural model we wish to link the empirical specification to (changing the notation from Section 6.1 slightly to accommodate our salience terms) is:

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

Letting $\mathbf{z_y}$ be the vector of year y prices, we linearize the specification around $\mathbf{z_1}$ using the following Taylor expansion:

$$\begin{aligned}
q_{i,2}(\mathbf{z_2}) &= f(\mathbf{z_2}) = q_{i,1} * exp\left(a + X_i * \delta + b_{i,ICR} * \left(\frac{log(P_{i,ICR,2})}{log(P_{i,ICR,1})}\right)\right) \\
&* \quad q_{i,1} * exp\left(b_{i,Donut} * \left(\frac{log(P_{i,Donut,2})}{log(P_{i,Donut,2})}\right)\right) \\
&* \quad q_{i,1} * exp(\theta_{Ded} * dedchg_i + \theta_{stark} * stark_i \\
&* \quad q_{i,1} * exp(u_{i,2}) \\
&\cong f(\mathbf{z_1}) + \left(\frac{\partial f(\mathbf{z_1})}{\partial P_{ICR}}\right)(P_{i,ICR,2} - P_{i,ICR,1}) + \left(\frac{\partial f(\mathbf{z_1})}{\partial P_{Don}}\right)(\mathbf{z_1})(P_{i,Don,2} - P_{i,Don,1}) \\
&+ \left(\frac{\partial f(\mathbf{z_1})}{\partial dedchg}\right) dedchg_i + \left(\frac{\partial f(\mathbf{z_1})}{\partial stark}\right) stark_i.
\end{aligned}$$

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

$$
\begin{aligned}
q_{i,2}(\mathbf{z_2}) &= q_{i,1} * exp(a + X_i * \delta + u_{i,2}) \\
&+ q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * \left(\frac{P_{i,ICR,2} - P_{i,ICR,1}}{P_{i,ICR,1}}\right) * b_{i,ICR} \\
&+ q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * \left(\frac{P_{i,Don,2} - P_{i,Don,1}}{P_{i,Don,1}}\right) * b_{i,Don} \\
&+ q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * (dedchg_i * \theta_{Ded} + stark_i * \theta_{Stark}) \\
&= q_{i,1} * exp(a + X_i * \delta + u_{i,2})(1 - b_{i,ICR} - b_{i,Don}) \\
&+ \left(\frac{q_{i,1}}{P_{i,ICR,1}}\right) * exp(a + X_i * \delta + u_{i,2}) * P_{i,ICR,2} * b_{i,ICR} \\
&+ \left(\frac{q_{i,1}}{P_{i,Don,1}}\right) * exp(a + X_i * \delta + u_{i,2}) * P_{i,Don,2} * b_{i,Don} \\
&+ q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * (dedchg_i * \theta_{Ded} + stark_i * \theta_{Stark})
\end{aligned}
$$

This gives us an expression for each of the coefficients from our structural model:

$$
\begin{aligned}
\alpha_i &= q_{i,1} * exp(a + X_i * \delta + u_{i,2})(1 - b_{i,ICR} - b_{i,Don}) \\
\eta_i * (\beta * \Pr(MP_i = ICR) + (1-\beta) * \Pr(CP_i = ICR)) &= \left(\frac{q_{i,1}}{P_{i,ICR,1}}\right) * exp(a + X_i * \delta + u_{i,2}) * b_{i,ICR} \\
\eta_i * (\beta * \Pr(MP_i = Don) + (1-\beta) * \Pr(CP_i = Don)) &= \left(\frac{q_{i,1}}{P_{i,Don,1}}\right) * exp(a + X_i * \delta + u_{i,2}) * b_{i,Don} \\
K_d &= q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * \theta_{Ded} \\
K_s &= q_{i,1} * exp(a + X_i * \delta + u_{i,2}) * \theta_{stark}
\end{aligned}
$$

These relations show that the linear structural constant α is highly individual-specific, varying with year 1 consumption and enrollee characteristics, so that exercise yields a *local* linearization of enrollees expected year 2 consumption relative to year 1 observed consumption and predicted year-to-year trends as a function of observables. 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.

In order to pool information across all regression sample individuals, 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. The results are in Table 13.

Table 13: Results of Quarterly ICR and Donut Price Regressions, with Stark Donut Coverage and Deductible Variables – High- and Low-Spending Enrollees, Pooled Years

		All Years Pooled (Pooled Regression)			
		Low-Spending Enrollee Response		High-Spending Enrollee Response	
Period	Price	Coef	SE	Coef	SE
Q1	ICR	-0.050	0.005 **	-0.018	0.007 *
Q1	Donut	0.028	0.007 **	-0.047	0.009 **
Q1	Ded. Chg	-0.046	0.006 **	-0.046	0.006 **
Q1	Stark	-0.012	0.007	-0.012	0.007
Q2	ICR	-0.045	0.005 **	-0.007	0.007
Q2	Donut	0.028	0.007 **	-0.055	0.012 **
Q2	Ded. Chg	-0.024	0.006 **	-0.024	0.006 **
Q2	Stark	-0.004	0.008	-0.004	0.008
Q3	ICR	-0.047	0.006 **	-0.005	0.009
Q3	Donut	0.012	0.008	-0.105	0.011 **
Q3	Ded. Chg	-0.019	0.006 **	-0.019	0.006 **
Q3	Stark	-0.003	0.009	-0.003	0.009
Q4	ICR	-0.060	0.007 **	0.026	0.010 **
Q4	Donut	0.003	0.009	-0.190	0.018 **
Q4	Ded. Chg	-0.013	0.007 *	-0.013	0.007 *
Q4	Stark	-0.046	0.009 **	-0.046	0.009 **

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. "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 elasticity information across groups of individuals whose price responses are expected to be similar by classifying sample individuals by quintile of $\left(\frac{q_{i,1}}{P_{i,1}}\right) * exp(a + X_i * \delta + u_{i,2})$, for each price P_i, giving us 25 bins overall – we estimate a single η_b for each bin b.[26] In order to pool dynamic information across individuals, 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, we use a generalized method of moments procedure to estimate our 25 static linear price response parameters η and our hyperbolic discount parameter β.

Table 14 displays the mean values for all model parameter estimates. The average linear price response $\eta = -1.7$ implies an average price elasticity of -0.09, which is similar to the literature. The quarterly hyperbolic discount factor $\beta = 0.19$ suggests substantial myopia.

Table 14: Estimated Structural Model Parameters, Pooled All Years

Description	Structural Parameter	Mean Estimate
Days supply (P=0)	α	374.349
Myopia	β	0.188
Marginal price effect	η	-1.681
Deductible effect	K_{deduct}	-8.773
Stark gap effect	K_{stark}	-6.588
Implied Elasticity	ε	-0.088

8 Discussion

In this project, we examine moral hazard in prescription drug consumption in the context of Medicare Part D, an insurance plan in which enrollees are exposed to substantial cost-

[26] The error term $u_{i,2}$ is not directly observed, so we use Duan's smearing technique to scale all transformed coefficients based on the distribution of the regression residuals.

sharing incentives and in which nonlinear, complex price schedules are found to lead to additional price effects beyond those anticipated by the designers.

Our identification strategy allows us to estimate static and dynamic price elasticities by leveraging variation in multiple budget set regions, focusing on groups whose marginal prices are in the interior of coverage phases. We rely on year-to-year changes in cost-sharing in multiple budget set phases for our identifying variation.

We demonstrate that, while enrollees' aggregate price responses are in line with many of those found in other health care contexts, Part D enrollees also exhibit certain behavioral biases in their consumption patterns due to the structure of Part D cost-sharing. In particular, we demonstrate evidence of imperfectly forward-looking behavior and of price salience effects.

The structural parameters implied by our linear estimates suggest that it matters not just what cost-sharing features of insurance plans *are*, but also *when* they are encountered in the year and *how* the prices are presented. A key policy implication of this research is that 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.

9 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.

Abaluck, J. and J. Gruber. 2013. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *National Bureau of Economic Research Working Paper No. 19163*. June.

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. 2012. "Moral Hazard in Health Insurance: How Important is Forward Looking Behavior?" *National Bureau of Economic Research Working Paper No. 17802*. February.

Blomquist, W. and W. Newey. 2002. "Nonparametric Estimation with Nonlinear Budget Sets." *Econometrica: Journal of the Econometric Society*. November. 70(6) 2455-2480.

Centers for Medicare and Medicaid Services (CMS). 2010. "Closing the Coverage Gap – Medicare Prescription Drugs Are Becoming More Affordable." *CMS Product No. 11493*. May.

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.

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. 2013. "The Response of Drug Expenditures to Non-Linear Contract Design: Evidence from Medicare Part D" *National Bureau of Economic Research Working Paper No. 19393*. August.

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.

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.

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 Drug Costs and Utilization." *National Bureau of Economic Research Working Paper No. 14326*. September.

Kling, J., S. Mullainathan, E. Shafir, L. Vermeulen, and M. Wrobel. 2008. "Misperception in Choosing Medicare Drug Plans." *Unpublished manuscript*. November.

Kowalski, T. 2011. "Economic Policy and the Financial Economic Crisis." *MPRA Working Papers, Faculty of International Business and Economics No. WP/2011/01*. August. 1-15.

Kowalski, A. 2012. "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*. May.

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 .

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.

Newhouse, J. and Rand Corporation. 1993. "Free For All? Lessons from the RAND Health Insurance Experiment." *Harvard University Press*. 450-475.

Newhouse, J. and Rand Corporation. 1993. "Free For All? Lessons from the RAND Health Insurance Experiment." *Harvard University Press*. 450-475.

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.