RAND Journal of Economics
Vol. 48, No. 4, Winter 2017
pp. 877–905

The impact of consumer inattention on insurer pricing in the Medicare Part D program

Kate Ho*

Joseph Hogan**

and

Fiona Scott Morton***

The Medicare Part D program relies on consumer choice to provide insurers with incentives to offer low-priced, high-quality pharmaceutical insurance plans. We demonstrate that consumers switch plans infrequently and search imperfectly. We estimate a model of consumer plan choice with inattentive consumers and show that high observed premiums are consistent with insurers profiting from consumer inertia. We estimate the reduction in steady state plan premiums if all consumers were attentive. An average consumer could save \$*1050 over three years; government savings in the same period could amount to* \$*1.3 billion or 1% of the cost of subsidizing the relevant enrollees.*

1. Introduction and motivation

■ The addition of pharmaceutical benefits to Medicare in 2006 was the largest expansion to the Medicare program since its inception. Not only is the program large, it is also innovative in design. Traditional Medicare Parts A and B are organized as a single-payer system; enrollees see the physician or hospital of their choice and Medicare pays a preset fee to that provider, leaving no role for an insurer. In contrast, Part D benefits are provided by private insurance companies that receive a subsidy from the government as well as payments from their enrollees. The legislation creates competition among plans for the business of enrollees, which is intended to drive drug prices and premiums to competitive levels. Each Medicare recipient can choose among the plans offered in her area based on monthly premiums, deductibles, plan formularies, out-of-pocket

* Columbia University and NBER; kh2214@columbia.edu.
** Columbia University; jph2154@columbia.edu.
*** Yale University and NBER; fiona.scottmorton@yale.edu.
We thank Mark Duggan, Liran Einav, Gautam Gowrisankaran, Ben Handel, Robin Lee, Bentley MacLeod, Aviv Nevo, Eric Johnson, and participants at numerous seminars and conferences for helpful comments. We especially thank Francesco Decarolis for sharing his data with us. All errors are our own.

costs (OOP costs or copayments) for drugs, and other factors, such as the brand of the insurer and the quality of customer service.

The premise of the Part D program was that consumers' ability to choose their preferred plan would discipline insurers into providing lower prices and higher quality than would be achieved through a government-run plan. Critically, these better outcomes require that consumers choose effectively, so that demand shifts to plans that consumers prefer because they offer low prices or high quality. If consumers choose plans randomly, a plan will have no incentive to lower its price because this will not affect enrollment, and markups will be high. In contrast, if many consumers choose to enroll in a plan that lowers its price, markups will fall as firms offer lower premiums to consumers.

This article compares Part D plan pricing when some consumers are inattentive to the case when all consumers are attentive. We demonstrate that, in reality, consumer choices are made with substantial frictions. Consumers rarely switch between plans and do not consistently shop for price and quality when they do switch. We provide evidence that, because of the absence of strong disciplining pressures from consumers, insurers set prices above the level they would choose if all consumers were attentive. Thus, plans extract high rents due to consumer inattention. Not only would improved consumer search benefit consumers directly, it would also lead to plan repricing that would save both consumers and the government significant sums. Our estimates indicate that removing inattention and allowing prices to adjust while leaving other sources of consumer preferences unchanged could reduce consumer expenditures by over \$1000 per enrollee over the three years 2007–2009. Under our assumptions, government program costs would fall by \$1.3 billion over the same period due to plan repricing. We find that the insurer response—lowering premiums—results in significant savings to both enrollees and taxpayers.

One concern when Part D began was that the prices the plans paid for drugs would rise because plans would lack the bargaining power of the government. Duggan and Scott Morton (2010) demonstrate that this did not happen. Rather, prices for treatments bought by the uninsured elderly fell by 20% when they joined Part D. Since the program's inception, increases in pharmaceutical prices have been restrained. This is due in part to aggressive use of generics by many insurers, but also to insurers' ability to bargain for rebates in exchange for favorable formulary placement and therefore market share. According to Congressional Budget Office (CBO) estimates, drug costs under the basic Part D benefit increased by only 1.8% per beneficiary from 2007–2010 net of rebates. The remainder of plan expenditures—approximately 20% of total costs according to the CBO—consists of administration, marketing, customer service, and like activities. The Personal Consumption Expenditures (PCE) deflator for services during this same time period increased at an average annual rate of 2.40%. Yet, despite these modest increases in the costs of providing a Part D plan, premiums in our data were on average 62.8% higher in 2009 than they were in 2006, the first year of the program, which corresponds to a 17.6% compound annual growth rate. The CBO estimates indicate that plan profits and administrative expenses per beneficiary (combined) grew at an average rage of 8.6% per year from 2007 to 2010.

These figures raise the question of why slow growth in the costs of drugs and plan administration were not passed back to consumers in the form of lower premiums. One possibility is that Part D may be well designed to create competition among treatments that keeps the prices of drugs low yet may not do so well at creating competition among plans in order to restrain the prices consumers face. Because the program is 75% subsidized by the federal government, any lack of effective competition would increase government expenditures as well as consumer costs. Our objective in this article is to investigate the extent to which consumer choice imperfections in this market impede competition between plans.

Section 2 of the article describes the Medicare Part D program and discusses reasons for search imperfections. In Sections 3 and 4, we review the literature related to consumer demand with choice frictions, in Medicare Part D and elsewhere, and the dynamics of pricing in that environment. In Section 5, we describe our data set, which provides detailed information on the choices and claims of nonsubsidized enrollees in New Jersey. In Section 6, we observe that consumers consistently make choices that are financially costly given their consumption patterns, and that this pattern of choosing expensive plans when cheaper ones are available does not appear to diminish with either experience in the program or time. Consumers seem to switch plans in response to "shocks" to their health or current plan characteristics, but are much less sensitive to changes in other plans. Motivated by these findings, we develop a two-stage consumer decision model for estimation, which accounts for inattention as a source of inertia. We identify the effect of consumer inattention separately from other potential sources of choice persistence, such as persistent heterogeneous unobserved preferences, using a detailed panel data set which documents the choices of new entrants to the Part D program and then follows each individual's choices over time. Our identification strategy is similar to that utilized in recent related articles that investigate the reasons for consumer choice persistence in other health insurance programs (Handel, 2013; Polyakova, 2014). The estimates indicate that inattention is an important part of the story.

Having established the behavior of consumers, we turn to analysis of the supply side of the Part D marketplace in Section 7. Using a data set of nationwide plan characteristics and enrollment, we show that plans with larger market shares set prices in a manner consistent with high choice frictions. We also document rapid growth in plan prices that is not accounted for by changes in costs and high dispersion in relatively homogeneous standard benefit plans that is indicative of search frictions.

The final section of the article, Section 8, uses the estimated demand model to conduct simulations that allow the supply side to adjust to a reduction in the proportion of consumers who are inattentive. We focus on plan premium choices. We abstract away from dynamics and model plans as profit-maximizing insurers that take into account the elasticity of demand, including consumers' attentiveness, when choosing a markup over cost. More attentive and price-elastic consumers will generate lower insurer margins. We use accounting data from the Part D program to estimate firm costs and then predict plans' optimal static premium bids under various assumptions regarding the proportion of consumers who are inattentive, and hence, the expected premium sensitivity faced by insurers. We show that the greater the percent inattentive, the higher the optimal premium bids chosen by plans. In our preferred simulation, removing inattention—and allowing plans to reprice in response to this change—would generate savings of $1050 per consumer over the years 2007–2009. In Section 9 we conclude with a discussion of the implications of our findings. The results indicate that even if consumers do not choose the lowest-cost plan for them, whether due to information processing costs or for other reasons, simply prompting them to choose a new plan every year has a substantial effect on costs through the channel of plan premiums. If we assume these changes can be generalized to plans nationwide, the federal government would save $1.3 billion between 2007–2009 if inattention was removed. Although we note that allowing for dynamic pricing would affect these predictions, perhaps generating an increase in premiums over time as plans "invest" to attract enrollees and then "harvest" their installed base, our static analyses are sufficient to demonstrate the substantial long-run savings to consumers and government that could result from increasing competition through reduced inattention.

Studies such as ours are crucial both to future policies concerning Part D plan design, information provision, and quality regulation, but also to those same issues in health insurance. Our results are relevant for policy and market design decisions in any healthcare market that relies on competition as a means to control costs and deliver quality.

2. Medicare Part D

■ Pharmaceutical benefits were not part of Medicare when it was first launched in 1965. However, the rising share of pharmaceuticals in the cost of healthcare created significant Out Of Pocket (OOP) expenditures for seniors and led to the creation of the Part D program under President George Bush in 2006. The novelty of this government benefit is the fact that it is essentially privatized: insurance companies and other sponsors compete to offer subsidized plans to enrollees. The sponsor is responsible for procuring the pharmaceutical treatments and administering the plan.

The Basic Part D plan is tightly regulated in its benefit levels so that there is limited flexibility for plans to reduce quality and thereby lower costs and attract enrollees. Plans must offer coverage at the standard benefit level, and each bid must be approved by the Centers for Medicare and Medicaid Services (CMS). The coverage rules include restrictions on plans' formularies, including which therapeutic categories or treatments must be covered. Plans are mandated to cover "all or substantially all" drugs within six "protected" drug treatment classes, as well as two or more drugs within roughly 150 smaller key formulary types. The protected classes include many treatments that would identify very sick patients, such as AIDS drugs, chemotherapy treatments, and antipsychotropics. Plans' placement of these drugs on their formulary is required, and the cost-sharing required of beneficiaries is carefully scrutinized by CMS to ensure plans are not discriminating against sick beneficiaries. Hence, it is not straightforward for a plan to avoid the sickest enrollees; this was particularly true in the first years of the program, when it was unclear which enrollees would have particular costs or utilization profiles and there was no usage history. Moreover, subsidy payments to plans are risk-adjusted according to their enrollees' demographics, and health status. There is an additional multiplier to increase the subsidy for low-income and institutionalized statuses. Thus, sponsors receive higher payments for sicker enrollees, which reduces their incentive to seek out healthy participants. In addition, plans must evaluate their predicted costs using CMS-specific actuarial models. This limits their ability to attract consumers by shifting costs to a part of the benefit that the enrollee has difficulty evaluating or will pay later. The result of this fairly tight regulatory environment is that the plan's premium emerges as its most salient characteristic for consumers, particularly for the Defined Standard Benefit plan.[1] We will see in our empirical work that consumers place high weight on a plan's premium when they make choices among plans. The deductible and other characteristics have an effect, but their empirical magnitude is much smaller than that of the premium.

Enrolling in Part D is voluntary, and one might be concerned that adverse selection would mean only sick seniors enroll. However, the subsidy for the program is set by legislation to be an average of 74.5% of costs, so for the vast majority of seniors enrolling is financially favorable (see Heiss, McFadden, and Winter, 2006) and most eligible seniors did enroll. In addition, the newly eligible who delay enrolling (perhaps until they become sick) are required to pay a higher price for coverage when they do join.

Many observers have noted that the Part D choice problem is difficult, and the empirical literature indicates that consumers do not choose plans that minimize their costs. In 2006, when the program began, there were at least 27 plans offered in each county in the United States. Enrollees had to consider how premiums varied across these plans, forecast their drug consumption in the year ahead, and compare the OOP costs for that set of drugs across plans. In addition, enrollees might receive an adverse health shock during the year that would change the set of medications demanded, necessitating the comparison of expected expenditures across plans. Furthermore, no major program like this existed in the United States at the time Part D began, so seniors likely had no experience attempting to make these calculations. Last, most Part D consumers are older Americans; outside the dual-eligible and disabled, Medicare eligibility begins at age 65. Finding a low-cost plan in the Part D program therefore requires the elderly to carry out a fairly difficult cognitive task.

Part D benefits are provided through two types of private insurance plans. The first is a simple prescription drug plan (PDP) which provides coverage only for prescription drug costs for seniors enrolled in the standard fee-for-service Medicare program (which does not cover drug costs). In 2006, 10.4 million people enrolled in PDPs. Medicare Advantage plans (MA-PD), for seniors who

[1] As we show in the article, enrollees can do better by searching for the plan-specific out-of-pocket payments for the particular drugs they will consume.

have opted out of standard Medicare, function similarly to a Health Maintenance Organization (HMO); such plans insure all Medicare-covered services, including hospital care and physician services as well as prescription drugs. In 2006, 5.5 million people enrolled in MA-PDs. By 2013, of the 32 million Part D enrollees, almost 20 million were enrolled in PDPs. MA-PD plans have particularly low enrollment in New Jersey, the state from which our data are taken: only 18–20% of New Jersey Part D enrollees were in MA-PD plans in 2006–2009, compared to 32–38% in the United States. overall. This article focuses solely on PDPs and prescription drug coverage. We assume that PDP enrollees do not consider enrolling in an MA-PD plan. We justify this assumption by noting both the low share of New Jersey MA-PD plans and the fact that moving from a stand-alone PDP to an MA-PD plan incurs the substantial cost of changing coverage (and potentially providers) for hospital and physician services as well as prescription drugs.

A fee-for-service Medicare enrollee can choose among all the PDPs offered in her region of the country. A plan sponsor contracts with CMS to offer a plan in one (or more) of the 34 defined regions of the United States. The actuarial value of the benefits offered by a plan must be at least as generous as those specified in the legislation. In the 2006 calendar year, this included a deductible of \$250, a 25% coinsurance rate for the next \$2000 in spending, no coverage for the next \$2850 (the “coverage gap”), and a 5% coinsurance rate in the “catastrophic region,” when OOP expenditures exceed \$3600. As these figures change annually, we report them through 2013 in Appendix Table 1 (available online). A sponsor may offer a basic plan with exactly this structure, or one that is actuarially equivalent—for example, with no deductible but higher cost-sharing. Enhanced plans have additional coverage beyond these levels and therefore higher expected costs and higher premiums.[2]

The way in which sponsors bid to participate in the program is important to an analysis of competition. Sponsors must apply to CMS with a bid at which each plan they wish to offer will provide the benefits of a basic plan to enrollees.[3] Importantly, the costs that the plan is meant to include in its bid are those it will expend to administer the plan, including, for example, the cost of drugs, overhead, and profit, and net of any costs paid by the enrollee, such as the deductible or copayments and reinsurance paid by CMS.[4] The bid is supposed to reflect the applicant’s estimate of its “average monthly revenue requirements” (i.e., how much it wants to be paid) to provide basic Part D benefits for a well-defined statistical person. CMS takes these bids and computes a “national average monthly bid amount” (NAMBA).[5] CMS uses the government subsidy percentage (74.5%) plus an estimate of its reinsurance costs and other payments to determine how much of the bid the beneficiaries must pay on average. This is called the beneficiary premium percentage, and in the first year of the program, it was 34%[6]. The Base Beneficiary Premium (BBP) is then the average bid (NAMBA) times the percentage payable by consumers. The premium for any given plan is this BBP adjusted by the full difference between the plan’s own bid and the NAMBA average. If a plan’s monthly bid is \$30 above NAMBA, then its premium will be \$30 above the BBP, and similarly if the bid is below the NAMBA (with the caveat that the premium is truncated at zero). This scheme makes the consumer bear higher premiums at the margin, which contributes to differences in premiums being important in consumer choice.

[2] The added benefit typically takes the form of either additional coverage in the coverage gap, reduced copayments, or coverage of certain drug types excluded from normal Part D coverage, such as cosmetic drugs and barbiturates. Plan sponsors offering plans with enhanced coverage must also offer a basic plan within the same region, and sponsors are prohibited from offering more than two enhanced plans in a given region. Enhanced plans do not receive higher subsidies, and any incremental costs are paid entirely by enrollees.

[3] Any costs of enhanced benefits in enhanced plans must be excluded at this stage.

[4] CMS may not bargain with plans over their bids. The agency may disallow a bid if some aspect of the plan, such as the formulary or the actuarial equivalence, does not follow regulations.

[5] In 2006, the various plans were equally weighted, but from 2008 onward, the NAMBA slowly transitioned to an enrollment-weighted average.

[6] The sum of the government subsidy and the beneficiary premium percentage is over 100% because part of the government subsidy is used for plan reinsurance rather than as a direct subsidy to premiums.

Two types of beneficiaries do not pay the full cost of Part D coverage. Approximately 6.3 million dual-eligible Medicaid recipients were automatically enrolled in Part D in 2006, as were an additional 2.2 million Low Income Subsidy (LIS) recipients. Premiums and OOP costs are fully paid by the government for the former, whereas the latter receive steep discounts. Foreseeing that LIS enrollees might be less price sensitive than regular enrollees, the Part D regulations only provide a full subsidy for LIS recipients who choose a plan with costs below the benchmark for their region.[7] If a plan loses benchmark status, its enrollees are automatically reassigned (equally across qualifying plans) to a benchmark plan in their region unless they choose to opt out and pay the cost difference themselves. Approximately 10% of enrollees in this category chose to opt out and become active "choosers" in 2007–2008 (Summer, Hoadley, and Hargrave, 2010). Although our demand model considers only non-LIS enrollees who are not dual-eligible, that article suggests that LIS enrollees who opted out behaved in a manner consistent with inattention, like the population we consider.[8]

The participation of private sector insurers in this new program in 2006 was voluntary and therefore uncertain. However, it turned out that many sponsors, both public and private, entered the Part D market in 2006. There were 1429 PDP plans offered nationwide in 2006 (though this had fallen to 1031 by 2013); every state had at least 27 PDPs every year during our sample period. Enrollees select one of these plans during the open enrollment period each November to take effect in the subsequent calendar year. The program includes many sources of aid for enrollees in making these decisions. Most importantly, CMS has created a website called "Planfinder" that allows a person to enter her zip code and any medications and see the plans in her area ranked according to OOP costs. The website also enables prospective enrollees to estimate costs in each plan under three health statuses (Poor/Good/Excellent), to estimate costs in standard benefit plans based on total expenditures in the previous year, and to filter plans based on premiums, deductibles, quality ratings, and brand names. A Medicare helpline connects the enrollee to a person who can use the Planfinder website on behalf of the caller in order to locate a good choice. However, conversations with CMS representatives suggest that very few enrollees make full use of the website. Pharmacies, community service centers, and other advocates offer advice. Survey evidence (Kaiser Family Foundation, 2006; Greenwald and West, 2007) indicates that enrollees rely on friends and family to help them choose a Part D plan, yet still find the choice process difficult.

3. Literature review: consumer demand

■ The introduction of Part D immediately created a literature evaluating outcomes from the novel program structure. An important early article documenting that the elderly do not choose optimally is that of Abaluck and Gruber (2011; hereafter, AG). Using a subset of claims data from 2005 and 2006 and a similar methodology to our own, the authors show that only 12% of consumers choose the lowest-cost plan; on average, consumers in their sample could save 30% of their Part D expenditure by switching to the cheapest plan. Consumers place a greater weight on premium than expected OOP costs, don't value risk reduction, and value certain plan characteristics well beyond the way those characteristics influence their measure of expected

[7] For the first three years of the program, the benchmark was calculated as the equal-weighted mean basic PDP plan premium in a region. In later years, it was an enrollment-weighted mean. Because lower cost plans have more enrollees, this policy change reduced the number of plans that qualified as benchmark over time.

[8] Because carriers set a single premium for a plan that enrolls both LIS and non-LIS consumers, there may be interactions between the two markets. For example, a strategy studied by Decarolis (2012) in the early years of Part D involved cycling of plans. A sponsor would raise the price of an existing plan above the benchmark, but introduce a new plan below the benchmark to catch auto-assigned LIS recipients, meanwhile, keeping any choosers and other enrollees in the original plan. We note that this cycling strategy was not used by all insurers, and declined over time. For example, Summer, Hoadley, and Hargrave (2010) report that by 2010, 92% of all auto-assignments were across corporate boundaries. In the analysis below, we find no evidence of this cycling behavior by plans in our New Jersey sample; we do not attempt to account for it in our model of supply.

costs. These results have been largely corroborated by Heiss et al. (2013) and Ketcham et al. (2012), among others.

Other studies have examined infrequent switching between plans as an explanation for inefficient consumer choice in the Part D market. In a field experiment, Kling et al. (2012) show that giving Part D consumers individualized information about which plans will generate the most cost savings for them can raise plan switching by 11% (from 17% to 28%) and move more people into low-cost plans. Ketcham, Lucarelli, and Powers (2015) use administrative data through 2010 to show that switching increases when more plans are available and that people become more responsive to large increases in their plans' costs over time. Polyakova (2016) estimates a model of plan choice featuring consumer switching costs and adverse selection, with unobservably riskier beneficiaries choosing more comprehensive coverage. She uses the model to simulate the effect of closing of the coverage gap on adverse selection and finds that switching costs inhibit the capacity of the regulation to eliminate sorting on risk. The presence of switching costs and consumer choice frictions has been documented in other health insurance markets by Handel (2013), among others.

Abaluck and Gruber followed up their results with a study of how enrollees' choices varied across the first four years of Part D (Abaluck and Gruber, 2013; hereafter, AG13). AG13 finds that consumers continue to make significant mistakes and that there is no measurable learning over time in their national sample. These findings are consistent with the estimates from our New Jersey sample. In both sets of results, consumers continue to be extremely sensitive to premiums. The empirical specification in AG13 is more reduced form than our model, but the two articles estimate similar levels of welfare loss from inertia. AG13 controls for brand fixed effects but still finds a strong role for inertia, concluding "rather than reflecting persistent unobserved factors of chosen plans, [inertia] reflect[s] either adjustment costs or inattention." Our article explores the inattention hypothesis in more detail. Our specification separately models consumer inattention, consumer valuation of the insurer's brand, and also persistent unobserved heterogeneity in preferences for a particular product. We continue to find an empirical role for inattention, even in this more sophisticated choice environment. AG13 concludes that choice inconsistencies are "driven by changes on the supply side that are not offset both because of inertia and because noninertial consumers still make inconsistent choices." By modelling the supply side, as we do in this article, we can simulate how insurers will set premiums in response to changing consumer attention. This step has received very little attention in the Part D literature. Ketcham, Kuminoff, and Powers (2016) predict the impact of various policies to reduce the impact of choice imperfections in Part D, for example, by reducing the size of the choice set, but they assume plan premium adjustments are designed to maintain the net revenue per enrollee that they earned prior to the policy. This assumption does not capture the impact of consumer inattention on plan markups that is the focus of this article.

There is a great deal of research in both psychology and economics literatures on consumer search and choice. Iyengar and Kamenica (2010) provide evidence that more options result in consumers making worse choices. In contrast to the prediction of a standard neoclassical model, more choice may not improve consumer welfare if it confuses consumers and leads them to seek simplicity. A large literature studies the importance of information processing costs to explain deviations from the choices expected of computationally unconstrained agents (see Sims, 2003; and Reis, 2006 for examples). Models of consumer search with learning, where each consumer uses the observed price of a single product to infer the prices likely to be set by other firms, also indicate that consumers may incur excessive costs by searching either too little or too much (e.g., Cabral and Fishman, 2012). Agarwal et al. (2010) show that the ability to make sound financial decisions declines with age. Because Part D enrollees are either disabled or elderly, and seem likely to experience cognitive costs of processing information, it may be reasonable to expect less optimal behavior from Part D consumers than from the population as a whole. These types of results have led some critics of Part D to call for CMS to limit the number of plans available to seniors. On the other hand, using data on private-sector health insurance, Dafny, Ho, and Varela

(2013) show that most employers offer very few choices to their employees and that employees would greatly value additional options. Moreover, the results from Stocking et al. (2014) suggest that merely limiting the number of available plans would not be sufficient, as this would limit competition and lead to higher prices. Thus, although the difficulty of choosing an insurance plan may lead consumers to choose expensive plans, it is not clear that limiting the range of options is the correct policy response.

Other authors have found evidence for inattention or lack of comparison shopping in complex and infrequent purchase decisions. In the auto insurance market, Honka (2014) finds that consumers face substantial switching costs, leading them to change plans infrequently, and that search costs lead those who switch to collect quotes from a relatively small number of insurers. Sallee (2014) uses the idea of rational inattention to explain why consumers underweight energy efficiency when purchasing durable goods. Busse, Simester, and Zettelmeyer (2010) find that consumers are inattentive and use a limited number of "cues" such as price promotions and mileage thresholds to evaluate auto purchases rather than actual prices and qualities. Luco (2016) and Illanes (2016) consider switching costs and firm competition in retirement investment choices. Hortaçsu, Madanizadeh and Puller (2015) examine consumer choices and switching behavior among retail electricity suppliers in Texas, and conclude that high search frictions lead to a high market share for the incumbent supplier.

4. Dynamics and pricing responses

■ Farrell and Klemperer (2007) surveys the substantial theoretical literature considering the effects of consumer switching costs and other sources of inertia on firm competition and pricing. There are two sets of results: the first relates to price changes over time whereas the second considers steady state price levels. We discuss both here but focus on the latter in our estimation.

Articles such as Klemperer (1987) and Klemperer (1995) argue that if firms cannot commit to future prices, consumer switching costs provide an incentive to "invest" and then "harvest." "Investing" is the process of building up market share through low prices in order to increase future profits, whereas "harvesting" is the process of reaping those profits by raising prices on an installed base. If the market begins in some particular period (as in 2006 for Medicare Part D), and all consumers have zero switching costs in that period, one might expect to see low initial prices and then price increases over time as the incentive to harvest the installed base increases relative to the incentive to attract new market entrants. In the longer term, once the market reaches steady state, multiple possible pricing patterns could emerge. If firms cannot discriminate between cohorts of consumers (as in the Part D application), new firms may choose a single price that is attractive to new consumers, and thereby effectively specialize in selling to new customers. Firms with old locked-in customers will choose a single price that is higher, and effectively specialize in selling only to old consumers, leading to cycling (Farrell and Klemperer, 2007). Alternatively, a firm may hold a "sale" in a particular period to attract new customers, whereas other firms pursue the same strategy in possibly different periods (Farrell and Shapiro, 1988; Padilla, 1995).

Several prior empirical studies have investigated the evidence for these dynamic pricing patterns in the presence of choice frictions. Ericson (2012) and Ericson (2014) analyze the insurer's problem in Medicare Part D, and show that firms initially set relatively low prices for newly introduced plans, but then raise prices as plans age, consistent with the "invest then harvest" dynamic. Similar questions have been studied empirically in other markets, for example, by Miller (2014) in the case of Medicare Advantage, and Cebul et al. (2011) in commercial health insurance. Decarolis (2015) and Decarolis et al. (2016) also study the supply side of the Part D market, paying particular attention to the interaction of low-income subsidy and other enrollees.

We show below that premium trends in the data are consistent with the predictions of these models: premiums increase substantially between 2006–2009, and they increase particularly for the plans with the greatest incentive to harvest their installed base (those with the greatest number of enrollees and in years with the smallest number of new entrants aging into Part D). However, we do not explicitly model price dynamics, and our counterfactual analyses do not predict price patterns over time. Instead, our simulations utilize the predictions of theoretical articles that analyze the impact of switching costs on price *levels*.

Beggs and Klemperer (1992) examine a no-sale equilibrium of an infinite-period duopoly model with consumer switching costs, in which in every period, new consumers arrive and a fraction of old consumers leaves. Firms cannot discriminate between these groups of consumers. Consumers are forward-looking and firms make dynamic profit-maximizing pricing decisions. Under the assumption that switching costs are sufficiently large that old consumers are locked into the product they have previously bought, there is a steady state Markov-Perfect Equilibrium where firm prices are higher than in the model without switching costs.[9] The intuition is that consumer lock-in gives the firm effective market power over some portion of consumers, which implies a price increase relative to the case with no switching costs. A similar intuition is provided in Radner (2003), which considers a model of "viscous demand," that is, where demand adjusts slowly to changes in prices. This viscosity provides the firm with a kind of market power, as it can raise its price above that of competitors without immediately losing all of its customers. In the homogeneous product duopoly case, there is a family of Nash equilibria where, once firms have achieved their target market shares and the total target market penetration, they both charge a price equal to consumers' willingness to pay (similar to a collusive outcome, with prices strictly greater than the equilibrium price without viscosity).

Farrell and Klemperer (2007) summarize these models, and other related articles, and conclude that there is a "strong presumption" that switching costs make markets less competitive, that is, lead to increased equilibrium prices.[10] Moreover, our setting has an additional feature that reinforces this conclusion. We find that consumer inertia in the Part D setting is caused by inattention (or asymmetric search costs) rather than switching costs of the conventional type. Because incumbent enrollees in a particular plan will only rarely notice other plans' prices, the incentive for firms to reduce prices in order to attract consumers from competitors ("invest") is small compared to a model with conventional switching costs. The intuition in Beggs and Klemperer (1992) is therefore likely to hold in the Part D context: equilibrium prices are likely to be higher than in the case without consumer inattention.[11]

Our counterfactual analyses investigate the magnitude of the steady state price increases likely to be generated by inattention, given our estimated model. We abstract from firm dynamic choices. We are interested in the steady state pricing of a marketplace of plans selling to fully attentive consumers compared to one pricing to inattentive consumers. We argue, following the intuition in Beggs and Klemperer (1992) and related articles, that the existence of inattentive enrollees has the effect of reducing the average elasticity faced by insurers. The true elasticity of any consumer's demand does not change, but a fraction of consumers is inattentive and therefore behaves inelastically. This group does not switch plans in response to a price increase, and therefore lowers the effective insurer elasticity of demand. We use this insight to generate a range of estimates of the premium effect of inattention.

[9] The authors note that the results will also hold if there is a startup cost K of trying a brand and K' of switching to a new brand, for K, K' sufficiently large.

[10] Work such as Dube, Hitsch, and Rossi (2010) shows that this result may not hold in cases where a firm's incumbent customers are not fully locked into a single firm. That article solves and/or simulates several simplified versions of a model with differentiated products, consumer switching costs, and imperfect lock-in. The authors show that as switching costs grow, equilibrium prices first fall and then increase, relative to the case with no switching costs. The intuition is that with low switching costs, the incentive for a firm to invest in future loyalty, and attract consumers from its competitors by lowering current prices can dominate the incentive to harvest. The competitor anticipates this and lowers its price to prevent the customer from switching. This effect is much less relevant for our setting because inattention is unlike other switching costs. Inattentive consumers do not notice a lower price and therefore cannot be attracted by it.

[11] The articles on consumer search and learning referenced above (e.g., Cabral and Fishman, 2012) also consider how firms price in response to consumer search. They contain similar intuition and make the point that the equilibrium outcome for prices depends on the size of the search cost relative to the variation in firm costs of production.

5. Data

■ Our primary data source, provided by the Centers for Medicare and Medicaid Services (CMS), contains information on prescriptions and plan choices for Part D enrollees from New Jersey in 2006–2009. Our data consist only of enrollees who did not have LIS status at any time and who were enrolled in stand-alone PDPs, rather than MA plans. Limiting the study to these enrollees reduced the population size from all New Jersey enrollees in PDP plans, of which there were between 527,000 and 545,000 from 2006 to 2009, to between 300,000 and 325,000 over the same time period. We chose New Jersey partially because it had a very low percentage of MA-PD enrollees compared to the national average—18%–20% of New Jersey enrollees were in MA-PD plans compared to a national average of 32%–38%—and because the total number of enrollees that met our criteria was not far above the CMS cutoff of 250,000. From this subpopulation, we drew a random sample in 2006 and a random sample of new enrollees in 2007–2009 to bring the total sample up to 250,000 enrollees. We limited the sample to unsubsidized PDP enrollees in order to focus on a setting where consumers had to pay the listed price for every plan and where plans had relatively standardized quality (not the case for MA-PD plans, which include medical as well as pharmacy benefits). Details of the data cleaning procedure are provided in the online Appendix.

Online appendix Table 2 shows the number of enrollees in our data set each year, ranging from 127,000 in the first year of the program up to 160,000 in 2009. Just over 60% of enrollees are female, and about 90% are white. The breakdown by age group is also shown in the table. Over our sample period the entering cohort, ages 65–69, grows in size from under 20% to almost 28% of the sample.[12] Because we have data from four years of the program, we can study the behavior of enrollees who have different numbers of years' experience in Part D. About 10% of each cohort leaves the program each year, and between 27,000 and 30,000 new enrollees enter each year.

The average quality of PDP plans nationally, as measured by the proportion of the 117 most-commonly prescribed chemical compounds covered by the plan, rises over time from 51% to 80%. Online appendix Table 3 summarizes the variation in this measure of quality across plans and over time. When weighted by enrollment, we see that consumers are slightly more likely to choose plans that include more drugs: the enrollment-weighted average coverage begins at 59% and rises to 82% by 2009. Our demand model accounts for this issue through consumers' expected out-of-pocket payments and through brand fixed effects and enhanced plan-year interactions.[13] Preferred pharmacy networks—which are not observed in our data—were not a significant factor during our time period. The Kaiser Family Foundation reports that only 6% of enrollees had a preferred pharmacy network in 2011, though they became popular shortly after that and expanded to 72% of enrollees by 2014.

For each enrollee, we estimate counterfactual costs in each plan (after discarding very small plans) holding consumption constant. Although Einav, Finkelstein, and Polyakova (2016) have shown that moral hazard affects an enrollee's drug consumption and, in addition, an enrollee might be elastic across therapeutic substitutes when she changes plans, dealing with these issues is beyond the scope of the current article. We follow the existing literature in our calculation of

[12] It may be that over time, employers and their about-to-be-retired employees no longer make other arrangements for pharmaceutical coverage, but build in to the employee benefit that he or she will use Part D. An evolution of this type would cause the flow rate into Part D at retirement to increase over time.

[13] One other dimension of quality that consumers might care about is customer service. CMS has a star rating system for enrollees to rate plans (with 3–5 stars available in each of 11–19 categories). Appendix Table 3 indicates that consumers may prefer higher-rated plans. However, the method used to assign star ratings changed dramatically between 2007 and 2008, making comparison between the 2006–2007 and 2008–2009 period difficult. There is evidence in prior articles that utilization management varies over time and across plans. The weighted average use of prior authorization for expensive drugs is 22% in 2014 (Hoadley, Summer, and Hargrave, 2014). We do not observe this in our data; it is captured in the brand and enhanced plan-by-year fixed effects in our utility equation.

counterfactual costs. Our methodology, described in detail in Section 2 of the Online appendix, combines elements of the techniques used in AG (2011) and Ketcham et al. (2012). First, we asked a physician to classify drugs as either chronic (taken regularly over a prolonged period) or acute (all other). We assume that chronic drug consumption is perfectly predicted by the patient and calculate the total underlying drug cost for each enrollee of the observed chronic drug prescriptions. For acute drugs, as in AG (2011), we assign each individual to a group of *ex ante* "similar" individuals, and assume that the consumer expects to incur a total per-month underlying drug cost equal to the median within her group. Following Ketcham et al. (2012), we then apply each plan's coverage terms (deductible, copayment, or coinsurance rate on each tier, gap coverage) to each individual and use his or her predicted total (chronic plus acute) monthly drug costs to predict total OOP spending, given these terms. This procedure yields estimates which closely track those we observe in the data for chosen plans. Although we expect there to be very little measurement error in the chronic OOP spending variable, as this is derived from observed utilization, there may be some measurement error in the acute OOP spending variable. Hence, in much of the analysis, we treat these variables separately.

6. The behavior of Part D enrollees

■ In this section, we explore the implications of the data for consumers' plan choice behavior. Further details and analyses are provided in Section 3 of the Appendix. Table 4 of that Appendix reports enrollee switching rates by demographic group in each of the observed open enrollment periods. From 2006–2007 a total of 19% of enrollees switch plans; this increases to 24% in 2007–2008 but falls to 8% in 2008–2009.[14] In every year, women and nonwhites are more likely to switch plans than other enrollees. The probability of switching increases monotonically with age. We create a group of those under-65 but eligible for Medicare due to disability. This group is similar in switching behavior to the 85+ group. The switching probability also decreases monotonically with income.[15]

We define the gap in payment as the expected OOP payment (including premium) in the chosen plan less the minimum expected OOP payment in any other plan in the choice set. We refer to this payment gap as "overspending" or gap spending. We note that, if consumers have preferences for nonprice characteristics, these may lead them to choose a plan other than the cheapest available; such a choice would not be an "error" and therefore, the term we use in the article is "overspending." Table 1 summarizes the level of overspending by year in our sample.[16]

In 2006, the first year of the program, the average amount paid above the minimum expected OOP payment available to the enrollee, including premium, was \$425.37, or 37% of the OOP payments. The percent and dollar amounts both fell in 2007 but then increased in both 2008 and 2009, to a level of \$436.96 or 36% of total spending in the final year of our sample. Thus, high spending is not declining over time in our sample. The data also indicate that part of the spending gap results from enrollees opting not to switch plans. Appendix Table 5 demonstrates that the spending gap is lower for consumers who have just switched plans, whereas it increases over time for nonswitchers. Appendix Table 6 shows that by 2009, over a quarter of switchers spent less than 110% of the cost of their estimated lowest-cost plan, whereas only 4% of those not switching achieved this.

One potential explanation for this behavior, which has been explored in numerous articles in this and other settings, is that consumers face switching costs, which lead to inertia. If switching costs were important, the consumers choosing to switch would be those for whom the value of switching was high enough to compensate them for these costs. Our data appear consistent with

[14] There are consumers who "passively" switch in the sense that the firm retires their plan and automatically moves them into a different plan run by the same firm, and we do not count these as switches.

[15] Income is measured as the median value in the enrollee's census tract; see the Appendix for details.

[16] We include both chronic and acute payments in our measure of OOP spending; the qualitatitive results change very little when we exclude acute spending.

TABLE 1 Overspending Relative to the Minimum Cost Plan By Part D Cohort

	Full Sample			New Enrollees			2006 Enrollees		
	Count	$ Error	% Error	Count	$ Error	% Error	Count	$ Error	% Error
2006	127,654	$425.37 ($369.50)	37.28 (22.38)	127,654	$425.37 ($369.49)	37.28 (22.38)	127,654	$425.37 ($369.49)	37.28 (22.38)
2007	141,897	$320.08 ($301.97)	29.61 (18.59)	28,460	$299.03 ($313.16)	30.12 (19.25)	113,437	$325.36 ($298.87)	29.48 (18.41)
2008	151,289	$378.72 ($348.80)	32.83 (17.98)	26,802	$331.88 ($346.83)	30.74 (18.91)	99,742	$387.50 ($346.24)	32.92 (17.49)
2009	159,906	$436.96 (359.44)	36.01 (16.49)	31,275	$371.78 ($371.34)	32.02 (18.44)	84,258	$459.19 ($353.25)	37.01 (15.61)

Notes: Predicted spending above the minimum by year. "%" is percent of enrollee's total OOP spending (including premium) in observed plan. Standard deviations in parentheses.

TABLE 2 Decomposition of Difference in Next-Year Overspending If Remain in Current Plan, Switchers versus Nonswitchers

Base Year	% from Change in Current Plan Prem.	% from Change in Current Plan TrOOP	% from Current Year	% from Change in Cheapest Plan Prem.	% from Change in Cheapest Plan TrOOP
2006	29.35%	−64.92%	173.89%	−16.77%	−21.54%
2007	71.76%	−0.62%	−9.98%	10.59%	28.26%
2008	57.11%	2.63%	2.28%	2.04%	35.93%

Notes: Decomposition of the difference between the change in overspending of switchers versus nonswitchers if they remain in their current plan. This difference is broken into five components: the contribution from the current year (defined as overspending in current year relative to lowest-cost plan), the increase in current plan premium and TrOOP, and the reduction in lowest-cost plan premium and TrOOP.

this idea. On average over all years and plans, switchers would overspend relative to the minimum-cost plan by $524 if they remained in their current plan, whereas the figure for nonswitchers is $338 on average.[17] We decompose this difference in next year's overspending between switchers (if they remained in the current plan) and nonswitchers, for each year 2006–2008, into five categories: overspending in the current year, the increase in the current plan's premium and in its predicted out-of-pocket cost (TrOOP) relative to the current year, and the reduction in the lowest-cost plan's premium and in its predicted TrOOP relative to the current year.[18] We report this decomposition, by base year, in Table 2, where a positive number indicates a larger contribution toward overspending for switchers than for nonswitchers. The decomposition is illuminating. Although the proportions differ over time, in two out of three years, over 55% of the difference between switchers' and nonswitchers' overspending if they remain in the current plan comes from changes in their current plan's premium.[19] In other words, a key distinguishing feature of switchers is not just that their value of switching plans is high, but that they also receive a signal of this fact in the form of a large increase in their current plan's premium.

Given these findings, we propose a slightly different explanation for the infrequent switching observed in the data. Rather than facing switching costs, consumers may be inattentive and in the absence of highly visible "prompts," may simply rollover their current plan choice. We argue in Section 4 of the Appendix that this behavior can be generated by a model where consumers have

[17] We exclude enrollees who enter or exit the program the following year from this analysis.

[18] Throughout the article, TrOOP refers to "true out-of-pocket costs," or OOP costs excluding premium, whereas OOP is the equivalent figure including premium.

[19] In the first year of the sample, the dominant factor is that switchers have larger errors in the current year than nonswitchers.

TABLE 3 Distribution of Shocks and Switching Likelihood

	Sample Distribution							
	No Acute Shock				Acute Shock			
	Neither	Premium Only	Coverage Only	Both	Neither	Premium Only	Coverage Only	Both
2006	50,503	49,954	3212	2824	3138	3082	170	554
2007	68,647	47,806	499	0	3488	4008	39	0
2008	78,980	37,081	5640	643	3651	2400	213	23

	Switching Likelihood							
	No Acute Shock				Acute Shock			
	Neither	Premium Only	Coverage Only	Both	Neither	Premium Only	Coverage Only	Both
2006	3.42%	32.68%	7.85%	40.16%	8.29%	54.77%	9.41%	42.24%
2007	7.34%	46.16%	0.20%	–	15.77%	57.53%	0.00%	–%
2008	1.72%	20.07%	14.52%	1.71%	2.63%	31.04%	13.62%	0.00%
Overall	4.10%	33.99%	11.46%	33.03%	8.82%	49.94%	10.66%	40.55%

Notes: Panel 1 sets out the number of enrollees with different types of shocks by year. Panel 2 summarizes switching probabilities by type of shock.

a cost of obtaining and processing information regarding alternative plan options and choose to incur this cost only when prompted by "cues" or "shocks" that are freely observed.

We investigate this hypothesis by considering three shocks to the consumer's own characteristics that could prompt her to incur the costs of search: two types of bad news concerning her current plan's characteristics for next year (the plan's premium will rise or coverage will fall noticeably) and an unusually high OOP payment driven by a health shock. We define a shock to premiums in the enrollee's current plan (v_p) as a premium increase of more than the weighted-median increase in the relevant year. A coverage shock (v_c) is defined as the plan dropping coverage in the coverage gap or moving from the Defined Standard Benefit to a different (tiered) system in the Pre-Initial Coverage Limit (Pre-ICL) phase. An enrollee is defined as having an acute shock (v_h) when she is in the top quintile of total drug cost as well as the top decile of either percent spending on acute drugs or deviation between predicted and observed spending. The distribution of these shocks in the population and their correlation with the decision to switch plans are shown in Table 3.[20] These three shocks appear to explain switching behavior well. Those who receive no shocks switch very infrequently, only 4% of the time, whereas those who receive multiple shocks are much more likely to switch plans.[21] Almost all switchers (87%) receive some shock in the year of the switch.

We present further evidence in support of consumer inattention in the Appendix. In particular, Appendix Table 7 sets out probit regressions of decision to switch plans on own-plan, low-cost plan, and personal characteristics. The estimates suggest that consumers' switching probabilities increase when their own plans' premiums and OOP costs rise, but we find no evidence that consumers respond to changes in premiums and costs for the lowest-cost plan available, the lowest-cost plan within-brand, or the average of the five lowest-cost plans.

□ **Consumer demand model.** We specify a simple two-stage model of consumer decision-making with inattention. We assume that each consumer i, once enrolled in a plan, ignores the

[20] The acute shock has a cross-year correlation of around .5, which is considerably lower than the cross-year correlation of other measures of sickness. Total spending, total supply, and acute supply each have a cross-year correlation between .8 and .9, implying that the acute shock is substantially less persistent than underlying health status.

[21] These findings are corroborated by Hoadley et al. (2013), who find that premium increases and removal of gap coverage are the best predictors of switching behavior.

choice problem until hit by a shock to the OOP costs of her current plan or to her health. We consider the same three shocks v_p, v_c, v_h defined above; these are assumed to have additively separable effects on her decision to reoptimize her plan choice. Additionally, the consumer could simply receive a random shock that causes awareness, for example, from a younger relative visiting the consumer and reviewing her plan choices. We label this shock v_e. The sum of these shocks creates a composite shock received by consumer i at time t:

$$v_{i,t} = v_{i,p,t}\beta_1 + v_{i,c,t}\beta_2 + v_{i,h,t}\beta_3 + v_{i,e,t}, \tag{1}$$

where the weights β allow the different shocks to have different effects on the propensity to search (e.g., shocks to premiums may increase the likelihood of switching more than other shocks).

When the composite shock $v_{i,t}$ is large enough, that is, when

$$v_{i,t} \geq \tilde{v}_{i,t}, \tag{2}$$

then the consumer becomes aware and decides to reoptimize her plan election. Here, $\tilde{v}_{i,t}$ is a function of consumer demographics related to health status and sensitivity to changes in plan characteristics: age groups, income quartiles, gender, and race. We also include year fixed effects in $\tilde{v}_{i,t}$ to account for differences in the environment (e.g., advertising, pharmacy, and government outreach) across our three different enrollment periods.

The second stage of the model examines how consumers who have decided to reoptimize choose whether to switch and to which plans. We assume that, once aware, consumer i makes a choice from the full choice set (including her current plan) based on the following utility from choosing plan j in year t:

$$\begin{aligned} u_{i,j,t} &= Tr\hat{O}OP_{i,j,t}\beta_1 + Premium_{j,t}[\beta_{2,1} + v_{i,p,t}\beta_{2,2}] + Ded_{j,t}\beta_{3,1} \\ &\quad + Gap_{j,t}[\beta_{4,1} + v_{i,c,t}\beta_{4,2} + v_{i,h,t}\beta_{4,3}] + X_{j,t}\beta_{5,i} + \epsilon_{i,j,t} \\ &= \delta_{i,j,t} + \epsilon_{i,j,t}, \end{aligned} \tag{3}$$

where expected chronic OOP spending excluding premium ($Tr\hat{O}OP_{i,j,t}$) is calculated using the method described above, $Premium_{j,t}$ and $Ded_{j,t}$ are annual premiums and deductibles, and $Gap_{j,t}$ is an indicator for any coverage in the gap. $X_{j,t}$ are nonprice plan characteristics, including brand fixed effects (defined at the carrier rather than the plan level) and an indicator for enhanced plans interacted with year fixed effects, and $\epsilon_{i,j,t}$ is an i.i.d. extreme value type 1 error term (assumed to be independent of $v_{i,e,t}$). We allow consumers prompted to search by shocks to premiums to place additional weight on premiums. Consumers experiencing shocks to coverage, or acute shocks, are permitted to place additional weight on the plan offering gap coverage.

We model persistent unobserved preference heterogeneity by including normally distributed random coefficients $\beta_{5,i}$ on fixed effects for the three dominant brands, which together have over 80% market share in 2006, and on the enhanced plan fixed effect. The model therefore allows choice persistence (such as a lack of switching away from a particular plan even when other plans reduce their premiums) to be caused either by heterogeneous preferences (some consumers have a very strong valuation for this brand that makes it worthwhile to remain enrolled even at a high relative price) or by inattention (consumers who are not affected by any of the previously defined shocks are unaware of other plan premium reductions).

The model is estimated using a random coefficients simulated maximum likelihood approach similar to that summarized in Train (2009). The likelihood function for each enrollee is predicted for a sequence of choices from entry into the Part D program until the end of our data panel. A full description of the empirical model is provided in the Appendix, where we also present estimates of the demand parameters. In all specifications, consistent with a model of inattention, the estimates indicate that consumers are significantly more likely to switch plans after receiving premium or coverage shocks or having an acute shock to their health. We now turn to the emphasis of the article, an analysis of firm behavior in the Part D marketplace.

TABLE 4 New Jersey Part D Market Summary Statistics

Year	Num. Plans	Enrollment	CR-4	HHI	Entering Plans	Enhanced Plans	Enhanced Mkt. Share	DSB Plans	DSB Mkt. Share
2006	44	281,128	0.862	0.259	44	17	12.27%	6	12.89%
2007	56	298,978	0.780	0.217	19	27	24.32%	8	10.49%
2008	57	304,198	0.617	0.157	9	29	28.62%	7	5.31%
2009	52	317,997	0.637	0.154	1	27	30.63%	5	0.48%
2010	46	329,178	0.660	0.163	2	24	30.43%	5	2.48%
2011	33	333,553	0.751	0.285	1	15	22.46%	4	2.53%
2012	30	343,886	0.753	0.281	3	14	24.00%	3	0.38%

Notes: Summary statistics on New Jersey Part D plans. “Num. Plans” is number of plans; “CR-4” is the four-firm concentration ratio; “HHI” is Hirschmann-Herfindahl Index; “Enhanced Mkt. Share” is the market share of enhanced plans; “DSB Plans” is the number of Defined Standard Benefit plans; and “DSB Mkt. Share” is the market share of Defined Standard Benefit plans. Source: Aggregate CMS data, generously provided by Francesco Decarolis. Total number of plans includes enhanced, Defined Standard Benefit (DSB), and other standard plans not following DSB coverage terms exactly. The latter are not listed separately in the table.

7. The supply side of the Part D market

□ **The New Jersey Part D market.** We begin with an overview of the supply side of the Part D market using a data set of Part D plans generously provided by Francesco Decarolis (Decarolis, 2015) that comprises CMS files on plans, ownership, enrollment, premiums, formularies, and other characteristics. It covers all plans in all regions of the United States for the years 2006–2012.[22] We focus on stand-alone Part D PDPs in New Jersey, as these are the plans which serve the consumers modelled in the previous section.

There were 44 PDP plans active in New Jersey in 2006, the first year of the Part D program; this is in line with an average of 42.2 plans per region nationwide. The New Jersey market is quite highly concentrated in every year of our data: measured in terms of enrollees, the four-firm concentration ratio begins at 0.862, declines to .617 in 2008, and rises again to .753 by 2012. Herfindahl indices show the same pattern. Our data agreement does not allow us to provide names for the large plans in our data. However, a table containing publicly available CMS information on the names and market shares of the five largest PDP plans in New Jersey in 2006, together with their brands, is provided in Appendix Table 13. There was little change in the rankings of these top five plans over the period of our data.[23]

There was some plan entry in New Jersey in the first several years of the program, but subsequent entry was limited. A total of 19 plans entered in 2007, joining 36 continuing from 2006, and 9 others entered in 2008, but from 2009 to 2012, no more than 3 plans entered in any year. After 2008 plan attrition reduced the number of active firms in every year from 57 down to 30 by 2012. Enhanced plans proliferated in the first few years of the program, going from 17 plans with a combined 12% market share in 2006 to 27 plans with a combined 31% market share in 2009. This coincided with a near-continuous shift away from Defined Standard Benefit plans; by 2012, only three such plans remained in the market, down from eight in 2007. These statistics, presented in Table 4, suggest an oligopolistic market characterized by increasing product differentiation and increasing concentration.

□ **Insurer pricing strategies.** We now consider the effect of consumer inattention, coupled with product differentiation and imperfect competition, on insurer pricing strategies in the Part D marketplace. We assume, as is traditional in industrial organization research, that insurers have rational expectations and are able to study the market in advance and choose an optimal strategy.

[22] See Decarolis (2015) for a detailed description of the data.

[23] The market shares listed in Table 4 and Appendix Table 13 are slightly different from the shares of the plans in the data used for our analysis, because as noted in the Appendix, we drop very small plans from our sample.

TABLE 5 Average Premium Increase and Percent of Plans with $10 Premium Increase

	Premium Increase				≥ $10 Premium Increase			
	Equal Basic	Equal Enhanced	Weighted Basic	Weighted Enhanced	Equal Basic	Equal Enhanced	Weighted Basic	Weighted Enhanced
2007	−$2.94	$1.01	−$2.20	$7.20	33.33%	40.74%	0.33%	10.53%
2008	$4.65	$11.50	$5.93	$14.45	39.29%	55.17%	24.10%	39.82%
2009	$6.20	$7.12	$3.68	$4.39	24.00%	33.33%	0.83%	39.31%
2010	$5.06	$1.77	$2.92	$5.44	21.74%	29.17%	1.19%	35.08%
2011	$1.04	$14.33	−$3.09	$2.84	11.11%	73.33%	6.50%	24.48%
2012	−$1.24	$6.52	$1.97	$2.02	12.50%	42.86%	0.16%	16.38%

Notes: Summary of premium changes ($ per enrollee per month) over time for New Jersey PDPs, by year and plan type.

We focus on the insurer's choice of premium. This is partly because the premium is the most important characteristic for consumer choice. It is also the metric CMS uses to approve plans and calculate each region's benchmark. Other aspects of the plan's strategy such as the design of the formulary (which, as noted, is quite tightly regulated) or gap coverage options are important areas for research (see, e.g., Carey, 2017; Einav, Finkelstein, and Polyakova, 2016; Lavetti and Simon, 2016) but are beyond the scope of the current article.

One would expect a profit-maximizing insurer to set its premiums in a way that took advantage of consumer choice frictions. In this section, we show that the patterns in the data are consistent with this intuition. Consider first, price dispersion. Varian (1980) features search in an environment of a homogeneous product, multiple sellers, and heterogeneous consumers. In this model, consumers do not engage in sequential search but rather "become informed" (perhaps by paying a cost) and at that point, know all prices. This model fits the situation where a consumer who has experienced a shock decides to reoptimize her plan choice, enters her Zip code and medications in the Part D website, and then has access to all firms and prices. The equilibrium symmetric outcome of Varian's model is price dispersion, which we certainly see in the Part D marketplace. In particular, Defined Standard Benefit plans are so tightly regulated as to represent a nearly homogeneous product. Each plan offers exactly the same financial tariff, and any given medicine is exactly the same in each plan. The plans differ by formulary, customer service, and brand. Different formularies will create differences in expected costs across individuals, but formularies are regulated by CMS to ensure that every therapeutic category has sufficient coverage and utilization management tools are appropriate—so the average value of each plan will be similar. Nevertheless, Table 14 in the Appendix shows that price dispersion persists among Defined Standard Benefit plans. Though the difference between the minimum and maximum premium is falling over time, there is still considerable variation in the cost of this close-to-homogeneous product by 2012.

As discussed in Section 4, consumer inertia is likely to have the effect of increasing equilibrium price levels and creating an upward slope to prices. The upward trend may not be a steady state phenomenon; it occurs because the entire market begins in 2006. Thus, every plan faces only elastic choosers in that year and no locked-in base. Table 5 shows that, consistent with these predictions, premiums increase on average almost every year from 2007–2012. The average annual premium increase for basic plans (weighted by enrollment) is small, less than $6 per month in every year. Premiums for enhanced plans increase more quickly; in 2008, the weighted-average premium increase for enhanced plans is over $14 per month, and in 2011 and 2012, smaller enhanced plans post large premium increases. The second panel of Table 5 flags plans that raise premiums by more than $10. For three years from 2008 to 2010, at least a third of enrollees in enhanced plans face large premium shocks, although the rate is lower in other years.

We can also use the intuition from the theory to predict differences in premium growth across insurers. First, the change in profit for a given change in price is a function of both the

TABLE 6 Estimated Coefficients from Regression on Annual Premium Increases ($)

	Model 1		Model 2		Model 3		Model 4	
	Coeff.	SE	Coeff.	SE	Coeff.	SE	Coeff.	SE
Lagged premium	−0.177***	0.008	−0.165***	0.008	−0.177***	0.008	−0.165**	0.008
Lagged # Tier 1 drugs	0.040***	0.005	0.037**	0.005	0.035**	0.005	0.031***	0.005
Lagged deductible	−0.009***	0.001	−0.008***	0.001	−0.009***	0.001	−0.007***	0.001
Lagged enhanced	1.448***	0.334	1.617***	0.335	1.442***	0.333	1.623***	0.334
Lagged gap coverage	5.773***	0.395	5.552***	0.396	5.750***	0.394	5.505***	0.396
Lagged market share	–	–	6.227***	1.220	–	–	6.716***	1.228
Enrollment growth rate	–	–	–	–	−3.288**	1.148	−4.011**	1.154
Brand Fixed Effect?	Yes		Yes		Yes		Yes	
Region Fixed Effect?	Yes		Yes		Yes		Yes	
N	7796		7796		7796		7796	
R^2	0.274		0.276		0.274		0.277	

Notes: Regression of premium increase (in $) on previous-year plan characteristics (national data). Enrollment growth rate is rate of growth for region's Part D program. Lagged market share is for this plan. "Coeff." is the estimated coefficient while "SE" is the standard error.

intensive margin (profit per enrollee) and the extensive margin (number of enrollees). Because larger firms have a larger intensive margin, we should expect large firms to raise prices more than smaller firms, all else equal. Second, we should expect slower premium growth when the number of consumers purchasing for the first time is high relative to the size of the installed base. Thus, premiums should rise more slowly in years with high attrition (e.g., high death rates) or large cohorts aging into the Part D program.[24] We estimate regressions of annual premium increases on lagged market shares, growth rates, and other plan variables that might affect costs for all PDP plans in the national data set.

Table 6 reports the results of the main specification. When we control for region and carrier fixed effects and coverage variables that may affect costs, lagged market shares significantly predict future increases in premiums, providing evidence in support of the first hypothesis. The estimates also indicate that the growth rate of enrollment in the region, which we treat as a proxy for new Part D enrollment, is negatively associated with price increases. This result provides evidence for the second hypothesis, that price increases should be small when there are relatively more unattached consumers to compete for. Taken together, the results of these regressions provide evidence consistent with the theory relating to price trends.

A further issue is that firms can sponsor more than one plan to offer more than one price. The work of Ericson (2012) and Decarolis (2015) leads us to investigate whether there is evidence of segmentation of consumers and price discrimination. In particular, the entry of basic "sister" plans may allow an existing plan to convert to enhanced status and raise its premium. The low-priced sister plan could enable the insurer to attract some enrollees who are auto assigned, or actively switch, to a low-priced plan. If carriers engage in this kind of consumer segmentation and "cycling," we should see higher premium growth of an existing plan when a new plan is added to the carrier's portfolio. The results of specifications including indicators for "sister" plan entry are provided in Appendix Table 15. We consider the impact of adding any "sister" plan and also the effect of adding a low-cost option: a plan whose premium is the lowest offered by the relevant carrier in the market. In both specifications, the relevant coefficient is negative and significant, implying that on average, plan premiums actually *fall* when a sister plan is introduced. These estimates suggest that, in contrast to Ericson (2012) and Decarolis (2015), in our sample,

[24] Because of our focus on shocks to consumers' attention and the dynamics of pricing, we do not estimate our motivating regression in levels like Polyakova (2016), but rather in premium changes. It is the increase in price that becomes more lucrative with an increase in installed base.

TABLE 7 Bids and Estimated Plan Costs for New Jersey PDP Plans

	Observed Bid	Observed Premium	Predicted Cost	Pred. Cost Net of TrOOP
2006	\$65.03 (\$26.68)	\$24.00 (\$10.23)	\$145.56 (\$39.10)	\$75.00 (\$26.23)
2007	\$64.93 (\$25.76)	\$25.05 (\$11.92)	\$162.24 (\$37.78)	\$84.86 (\$19.01)
2008	\$92.28 (\$31.04)	\$35.29 (\$15.83)	\$153.18 (\$43.43)	\$85.60 (\$33.00)
2009	\$100.97 (\$29.90)	\$40.34 (\$15.22)	\$154.03 (\$40.69)	\$87.90 (\$40.53)

Notes: Summary of weighted-average observed bids, observed premiums, predicted costs to the plan, and predicted costs net of enrollee out-of-pocket payments. All figures are per enrollee per month. Weighted standard deviations in parentheses; weighted by enrollment.

premiums do not increase more than average when the carrier adds a new plan to the portfolio. We will not model this cycling behavior in the simulations below.

□ **Insurer cost estimates.** Our next step is to use accounting data (our claims data from New Jersey) to estimate each plan's average cost per enrollee. These costs will be used as an input to the counterfactual premium simulations in the following section.

The claims data indicate the gross drug cost for every claim, including the drug ingredient cost plus the dispensing fee and sales tax paid to the pharmacy, but not accounting for manufacturer rebates or plan administrative costs. For each branded drug, we find the average gross drug cost of a 30-day supply across all plans and all encounters in the relevant year and apply a 20% rebate to that average cost. For generic drugs, we assume a \$4 cost per 30-day supply for all plans.[25] We use these figures, and the observed drug utilization for each enrollee, to predict an average drug cost net of rebates per enrollee per year. Our methodology also accounts for the fact that, as part of its risk-adjustment strategy, the government covers 80% of all drug costs in the catastrophic phase, so that the plan pays at most 20% of these costs.[26] We reduce the effect of outliers by winsorizing the estimated per-person costs at the 2.5% level (i.e., replacing the top and bottom 2.5% with the 2.5th and 97.5th percentile, respectively). We then compute the average per-person drug cost of the plan's beneficiaries.

We expect these cost estimates to be fairly accurate for several reasons. The cost of generics and average branded discounts are unlikely to vary across plans as they might in a less regulated market. Because formularies in Part D are highly regulated (e.g., the protected classes must contain substantially all drugs in the class and other classes must have a certain number of drugs covered) and may not be designed to drive away sick enrollees, plans cannot differ substantially in the "tightness" of their formularies—and this tightness is what determines elasticity and therefore the size of a brand rebate (Duggan and Scott Morton, 2010). Furthermore, we know exactly which drugs, and how many units of each, were dispensed by each plan.

Finally, we need to account for plan administrative costs. Sullivan (2013) notes that the National Health Expenditure Accounts (NHEA) includes the administrative costs of Medicare Advantage plans and Part D plans in its report of total Medicare administrative costs. We use this fact, and data from the NHEA for 2006–2010, to back out administrative expenses of 14%–16% of total costs—or 16%–19% of nonadministrative costs—for Parts C and D combined. We therefore inflate the estimated average plan-level drug cost per enrollee per year by 120% to account for administrative costs. Additional details on the construction of the cost estimates is provided in Section 5 of the Appendix.

The resulting plan costs per enrollee are summarized in Table 7. We report weighted averages and standard deviations of both the total cost per enrollee and the estimated cost net of enrollee

[25] A study by the Department of Health and Human Services Inspector General (Levinson, 2011) found that, in 2009, rebates reduced Part D drug expenditures by 19% on average for the 100 highest-volume brand-name drugs. Our assumption regarding generic drug costs is based on Walmart's well-known "\$4 for any generic prescription" program.

[26] In most cases, the beneficary pays a 5% copay in the catastrophic phase, so for branded drug events, we assume the plan pays 15%. Fewer than 5% of enrollees reach this phase, so this has little effect on predicted plan costs.

out-of-pocket payments.[27] The latter will be the cost variable used as an input into the premium-setting simulations below. Finally, we report for comparison the weighted-average observed bid and observed premium separately for each year of our data. Observed bids are about \$10 lower than predicted costs net of TrOOP on average in 2006, the first year of the program. Observed bids fall slightly in the second year, and this together with an increase in estimated costs implies a lower average markup. Bids increase much faster than predicted costs in the following two years.

The plan markup does not equal the bid less the cost, and for this reason we do not report a markup estimate based on these data. Plan revenues also include an additional premium amount for enhanced plans plus reinsurance payments from CMS. The plan profit equation in Section 8 provides more details of these elements of revenue. For now, we note that the estimates in Table 7 clearly indicate that plan margins did not converge toward zero over the first few years of the program.

8. Counterfactual simulations

■ Previous studies have considered the effects of various interventions designed to ease the decision-making process at fixed prices. For example, in a randomized experiment, Kling et al. (2012) provide information to Part D enrollees regarding their best plan choice and find that it increases the probability of switching by 11 percentage points. Abaluck and Gruber (2013) predict that if an intervention could make consumers fully informed and fully rational, they would choose plans that reduced their costs by about 27%. However, these articles do not account for plans *repricing* in response to changes in consumer behavior, potentially further lowering program costs. In this section, we use the estimated demand model and our measure of firm costs per enrollee as inputs to counterfactual analyses. We calculate how insurer markups respond to consumer behavior.

We investigate the magnitude of the steady state price increases likely to be generated by inattention given our estimated model. Note first that although the firm pricing problem in the observed data is dynamic, the dynamics come from the opening of the market and from the forward-looking choices of firms who anticipate future consumers will be "sticky," so that prices chosen in one period affect enrollment in future periods. We abstract away from these dynamic pricing issues in the simulation in order to compare steady states. We argue that the change in steady state outcomes is of most policy relevance for a program that is long-lived. We assume in our demand model that attentive consumers are not forward-looking or strategic; rather, they choose the plan they most prefer today based on today's characteristics. (This seems very likely, given what we know about the sophistication of these consumers and their uncertainty about how the market will develop in the future.) This estimation therefore yields preference parameters that are valid representations of consumer behavior, regardless of the strategy the firms may be pursuing. We can then use these preference parameters combined with our cost data to generate a variety of different static equilibrium markups, corresponding to different levels of inattention and hence different premium elasticities, using a simple system of static first-order conditions. We consider only non-LIS enrollees and focus on the simple situation where consumer preferences are not affected by shocks experienced in the previous year.

We focus on the premium markup for several reasons. First, the results of Abaluck and Gruber and others indicate that consumers choose poorly in this marketplace and are unable to properly weight attributes of a plan such as out-of-pocket costs and deductibles. Consumers often choose plans based largely on the premium. Our demand estimates are consistent with this finding, implying that removing inattention (without addressing other choice frictions) would affect consumer sensitivity to premiums much more than their response to out-of-pocket costs. That is, we focus on the plan characteristic that is most likely to respond to the removal of inattention. Further, basic plans are constrained to offer a tariff that is actuarially equivalent to

[27] We truncate the plan-level average cost net of OOP payments at zero; this step involves only a few plans.

that set out in the law, so they are financially fairly homogeneous to the average consumer except for the premium. Although insurers may design the formulary, and hence out-of-pocket payments, of their plans (particularly enhanced plans) in a strategic way to influence enrollee choices and utilization, the plan's bid is the choice over which it has most discretion and the one that feeds directly into the premium calculation, and therefore its competitiveness in the marketplace. We focus on this choice, leaving issues related to formulary design for future research.[28]

We begin by predicting plans' optimal premiums in the simple scenario of full attention; details of our method are provided in the following subsection. The next step is to use the premium, coverage, and acute health shocks observed in the data (generated by realized plan premium and coverage changes) to calculate the proportion of consumers that our model (equation (2)) predicts will be attentive in each year. For example, the shocks in the data imply that 37% of consumers will reoptimize in 2007, whereas 100% are assumed to optimize in 2006. We assume that attentive consumers make choices to maximize the estimated utility equation (3) whereas inattentive consumers have a zero price coefficient in that equation. This provides a simple way to translate the "probability of attention" implied by the mix of attentive and inattentive consumers in the data into an expected premium sensitivity that is faced by insurers.[29] It is then straightforward to repeat the exercise of predicting equilibrium premiums using static first-order conditions given the new (less negative) premium coefficient.

This method abstracts away from several complicating factors. First, as noted, our desired simulation does not require us to address dynamics at all. Second, we do not compare our simulated outcomes to data because the data contain in them the impact of the dynamics—which might lower the absolute level of markups in early years, for example—and therefore are not directly comparable to our predictions. Our objective is to provide an approximate magnitude of the impact of consumer inattention on the steady state markup. We therefore calculate a range of possible premium effects under different assumptions regarding the mix of attentive and inattentive consumers in the population. As we vary the proportion of consumers who are attentive, we vary the price elasticity of demand facing the insurer, and hence its incentive to choose a particular premium level. We assess the range of likely premium effects of inattention by reporting predicted consumer spending when the percent attentive varies between 20% and 80% of the population.

A second caveat is that we do not model each individual plan at the micro level to make the link between that plan's premium choice and the attentiveness of its own incumbent enrollees. That is, our analysis does not model the strategy of each plan to raise premiums and lose some (newly attentive) consumers. We instead assume a population-level probability of consumer attentiveness, and then calculate premiums that generate an outcome at the population level. Our preferred method uses an iterative algorithm to address the issue that the premiums chosen by particular plans may generate premium increases (and hence, shocks) that affect the proportion of enrollees choosing to reoptimize, implying that assuming a fixed probability of attentiveness to which plans respond when choosing premiums may generate predictions that are internally inconsistent. Our iterative algorithm proceeds as follows. The predicted premium choices in iteration k are used to generate predicted premium shocks and therefore a probability of attentiveness (from equation (2)) and new premium coefficient (in equation (3)). That premium coefficient is used to predict new premium choices in iteration $k + 1$. We repeat until convergence, thereby ensuring that the population percent attentive is consistent with the proportion of plans generating premium shocks.

We note that some price changes over time are predicted in our simulations, despite their being a steady state exercise. We do not mean these to be interpreted as "dynamics." Rather, they are changes in static predictions that occur due to changing inputs, such as underlying plan

[28] Articles such as Carey (2017), Einav, Finkelstein, and Polyakova (2016), and Lavetti and Simon (2016) consider strategic factors affecting insurers' formulary choices in Part D.

[29] We generate a new utility equation that is the same as equation (3), except that the premium coefficient is scaled by the percent of consumers who are attentive. It would be more accurate to assume that inattentive consumers had a zero price elasticity, but this would substantially complicate the simulations.

costs, and to the fact that the proportion of enrollees new to Part D differs by year. For example, premiums are predicted to be low in 2006 in every simulation, as all consumers are new to the market in that first year and therefore are assumed to be attentive.

We now discuss the possible sources of error in our approach. As noted above, we expect our approximation of costs to be fairly accurate. Any mismeasurement would in any case be likely to affect cost levels rather than the differences across our simulations. Second, our model does not account for the possibility that consumers may in fact be forward-looking, choosing plans with relatively high prices because they expect those plans to have lower prices in the future when they are inattentive. However, there is a large amount of research into the behavior of Part D consumers, none of which indicates that this kind of forward-looking planning is probable. Third, there might be a steady state that is a perpetual cycle (perhaps with entry and price changes, or LIS cycling as in Decarolis, 2015) and not the static outcome we simulate. Our approach also abstracts away from other reasons why premiums might change over time, including insurer learning and low pricing in early years to attract enrollees with the goal of switching them to MA plans. Although these dynamic effects are certainly possible, we showed in Section 7 that we found no evidence of premium cycling in response to sister plan entry in our sample. Further, we argue that short-term issues such as plan learning are of second-order importance for the steady state premium effect of inattention that is our focus.

An alternative approach to our counterfactuals would have been to simulate the dynamic path of premiums in the baseline model with inattention. However, predicting the equilibrium of a dynamic pricing game with many firms is difficult. Articles on the methodological frontier have either considered very simple markets with two firms and small numbers of consumer types (e.g., Dube, Hitsch, and Rossi, 2010) or made other simplifying assumptions, for example, of a finite time horizon, or the simplification that markets are large enough that the random evolution of individual firms "averages out" and each firm can be assumed to respond to a long-run average industry state rather than the predicted current choices of its competitors (Weintraub, Benkard, and Van Roy, 2008; and applications such as Miller, 2014). None of these assumptions seems reasonable in our setting. We argue that our simulations provide a reasonable first approximation to the substantial cost savings, for both consumers and the federal government, that could arise from policies to increase consumer attentiveness. We leave the specification and estimation of a dynamic premium-setting model as an important path for future research.

□ **Counterfactual details.** We simulate consumer and plan choices using the fixed sample of 40 New Jersey PDP plans that entered the market in 2006; we follow each plan through to its exit from the market.[30] By limiting ourselves to these 40 plans, we ensure that the reported numbers focus only on within-plan price and spending changes.

Recall that each insurer submits a bid for each plan. That bid determines the price consumers face (the amount over the base beneficiary premium). This institutional reality requires us to reframe the static price-setting game that is standard in the industrial organization literature as a game where insurers choose bids, given a prediction of the implications for the premiums faced by consumers. Importantly, each basic plan must offer actuarially equivalent coverage if it does not follow the tariff set out by law. This means that for a statistical person, the mean of OOP charges must be the same in expectation for all basic plans, so plans cannot respond to increased consumer premium sensitivity by reducing premiums while increasing average OOP charges. Additionally, the subsidy for each enrollee is risk-adjusted depending on age, chronic conditions, LIS, and institutional status. Although the risk-adjustment mechanism is potentially manipulable, risk-adjusted subsidies plus the high share of catastrophic costs paid by CMS (80%) mean it will be difficult for firms to immediately determine whether particular enrollees are profitable or not, and the computation will be complex. We therefore abstract from selection issues as we model

[30] There were actually 44 PDP plans in New Jersey in 2006 (Table 4); as noted in the Appendix, we drop the smallest plans in the sample, so we focus on 40 of these 44. 31 of the 40 plans were still active in 2009.

the behavior of insurers. We model insurers' choices of bids while holding the schedule of OOP charges fixed.

We write plan j's variable profit in year t as:

$$\pi_{j,t} = (B_{j,t} + E_{j,t} - C_{j,t})N_{j,t}, \tag{4}$$

where $B_{j,t}$ is the bid made to CMS reflecting the plan's average monthly revenue requirement per enrollee in a basic plan (including profit), $E_{j,t}$ is the additonal amount charged to enrollees in an enhanced plan (the "enhanced premium"; this is zero when j is a basic plan), $C_{j,t}$ is the plan's cost per enrollee net of enrollee OOP payments, and $N_{j,t}$ is its number of enrollees.

The premium charged to enrollees in a basic plan is the difference between the bid and the proportion of the NAMBA that is subsidized by the government:

$$Premium^{Basic}_{j,t} = B_{j,t} - \gamma_t NAMBA_t = \left(1 - \frac{\gamma_t}{J_t}\right) B_{j,t} - \frac{\gamma_t}{J_t}\Sigma_{k\neq j} B_{k,t}, \tag{5}$$

where γ_t is the proportion of the NAMBA that is paid by the government and J_t is the number of Part D plans included in the average in year t.[31] This expression reflects the fact that, in the first two years of the program, the NAMBA was an unweighted national average of bids for all MA and PDP plans. From 2008 on, CMS phased in the implementation of a weighted average, where the weight was the plan's enrollment.[32]

We take several steps to account for CMS's risk-adjustment strategy. The government subsidy, which is written into law at 74.5% of the NAMBA, is split between a premium subsidy and reinsurance or risk-adjustment payments. The latter include a commitment to pay 80% of the total cost of drugs above each enrollee's catastrophic threshold and payments to keep plans within symmetric risk corridors that limit their overall losses and profits. We adjust our measure of plan costs per enrollee to take account of the catastrophic drug subsidies as described in the previous section. We use the true proportion of the NAMBA that is paid by the government in every year (which is observed in our data, e.g., 66% in 2006) as an input to the premium calculation in equation (5). We assume that the remaining risk-adjustment payments offset the effect of enrollee selection on plan costs, that is, the cost per enrollee does not change with enrollees' plan choices in our simulations.

Under each counterfactual scenario we consider a single-stage consumer demand system. We use the estimated parameters of the utility equation described in Section 6 (the full specification, model 4 of Appendix Table 11), but set the coefficients on premium, coverage, and acute health shocks to zero. The premium coefficient is adjusted as described above to approximate the effect of a particular mix of attentive and inattentive consumers on premium responsiveness. The resulting utility equation can be written as:

$$\begin{aligned} u_{i,j,t} &= Tr\hat{O}OP_{i,j,t}\beta_1 + Premium_{j,t}\beta_{2,1} + Ded_{j,t}\beta_{3,1} + Gap_{j,t}\beta_{4,1} + X_{j,t}\beta_{5,i} + \epsilon_{i,j,t} \\ &= \lambda_{i,j,t}(\beta_{5,i}) + Premium_{j,t}\beta_{2,1} + \epsilon_{i,j,t} \\ &= \delta_{i,j,t}(\beta_{5,i}) + \epsilon_{i,j,t}, \end{aligned} \tag{6}$$

where $Premium_{j,t}$ includes the enhanced premium where relevant and $\beta_{2,1}$ is adjusted across simulations. $\lambda_{i,j,t}(.)$ includes all consumer and plan-specific variables in the estimated utility equation except the premium; it is a function of $\beta_{5,i}$, the random coefficients on the three dominant

[31] CMS requires that the basic premium never fall below zero. Our simulations account for this constraint. However, we note that the constraint is not binding for PDPs in our data, although MA-PDs, which bundle prescription drug insurance with Medicare Part C insurance and whose bids are included in the NAMBA, often have very low premium bids.

[32] The premium charged to enhanced plan enrollees is the basic premium defined in equation (5) plus the enhanced premium $E_{j,t}$. The enhanced premium is negotiated between the carrier and CMS and is meant to comprise the average additional cost of enhanced benefits provided to enrollees in the plan. It is not subsidized by CMS. We observe this variable in the data for every plan-year and account for it in our simulations under the assumption that it does not change in response to changes in enrollee attention.

brands and the interactions between enhanced plan and year fixed effects. This utility equation can be used to predict plan enrollment $N_{j,t}$ under any set of plan characteristics:

$$N_{j,t} = \sum_{i=1}^{N_t} \int_{\beta_{5,i}} \frac{e^{\delta_{i,j,t}(\beta_{5,i})}}{\sum_{k=1}^{J_t} e^{\delta_{i,k,t}(\beta_{5,i})}} \partial F(\beta_{5,i})$$

$$= \sum_{i=1}^{N_t} \int_{\beta_{5,i}} \Lambda_{i,j,t}(\lambda_{i,j,t}(\beta_{5,i}), \lambda_{i,-j,t}(\beta_{5,i}), Premium_{j,t}, Premium_{-j,t}) \, \partial F(\beta_{5,i}). \quad (7)$$

Here, $\Lambda_{i,j,t}(.)$ is the predicted probability that consumer i chooses plan j in period t; it is a function of all plan characteristics, including their premiums. We consider plans' optimal choices in the static bid-setting game. The first-order condition for plan profits with respect to the bid $B_{j,t}$ is:

$$(B_{j,t} + E_{j,t} - C_{j,t}) \frac{\partial N_{j,t}}{\partial B_{j,t}} + N_{j,t} = 0. \quad (8)$$

Calculating the derivative $\frac{\partial N_{j,t}}{\partial B_{j,t}}$ requires us to predict the effect of a change in the bid $B_{j,t}$ on the premium. We use the expression in equation (5) under the assumption that the NAMBA is an (unweighted) national average for MA-PD and PDP plans and that plans internalize their impact on the NAMBA, and therefore on the government subsidy, when choosing their bids.[33] We predict the resulting effect on enrollment using equation (7). The first-order condition simplifies to:

$$N_{j,t} + (B_{j,t} + E_{j,t} - C_{j,t}) \left\{ \Sigma_{i=1}^{N_t} \beta_{2,1} \left[\int_{\beta_{5,i}} \Lambda_{i,j,t}(.)(1 - \Lambda_{i,j,t}(.)) \, \partial F(\beta_{5,i}) \right] \frac{J_t - \gamma_t}{J_t} \right.$$
$$\left. + \Sigma_{k \neq j} \beta_{2,1} \left[\int_{\beta_{5,i}} \Lambda_{i,j,t}(.) \Lambda_{i,k,t}(.) \, \partial F(\beta_{5,i}) \right] \frac{\gamma_t}{J_t} \right\} = 0,$$

where we omit the arguments of $\Lambda_{i,j,t}(.)$ for ease of exposition. All plans' bids enter this equation through $\Lambda_{i,j,t}(.)$ as well as through $B_{j,t}$. We solve this system of equations to obtain the implied new equilibrium for bids. Additional details of this derivation are provided in Section 5 of the Appendix.

□ **Counterfactual Results.** Tables 8–10 report our results. Table 8 sets out the cross-enrollee average premium costs and OOP spending (including premiums) predicted by our model with and without inattention. The first column shows the predictions under full attention. Column 2 ("Inattention (Using Observed Shocks)") shows the same simulated costs when the shocks observed in the data between 2006–2008 are used to predict the proportion of consumers who are attentive the following year (given equation (2)), and the premium coefficient in equation (6) is scaled down accordingly. The proportion attentive is reported in the column labelled "P(attn.)." It equals 1 in 2006, the first year of the program; predicted spending with and without inattention are equal in that year. In the subsequent three years, the percent attentive varies between 0.33–0.45. Premiums and OOP spending are lower under full attention in each of these years for two reasons. First, plans reduce their premium bids in response to the higher premium elasticity, generating a reduction in premiums charged. Second, enrollees place a higher (negative) weight on premiums in the utility equation, and hence choose lower-premium plans. These results indicate a large saving from the simulated change in elasticity due to attention: a total out-of-pocket cost saving (including premiums) of $1154.20 per enrollee over three years, or 25.6% of total spending.

[33] We account for the fact that a change in one plan's bid will affect all plans' premiums via the subsidy. We use national NAMBA figures published in annual press releases as an input to this analysis. The bid-setting game is solved for New Jersey PDP plans, assuming that PDP plans outside New Jersey will change their bids proportionately with New Jersey PDP plans, but holding fixed the bids of MA plans both within and outside this state. See the Appendix for details.

TABLE 8 Simulated per-Person Spending with Premium Adjustments

	Full Attention		Inattention (Using Observed Shocks)			Inattention (Iterative Algorithm)		
	Premium	OOP	Premium	OOP	P(attn.)	Premium	OOP	P(attn.)
2006	$190.19	$1111.90	$190.19	$1111.90	1	$190.19	$1111.90	1
2007	$263.62	$1149.00	$678.26	$1545.80	0.37	$611.49	$1480.30	0.42
2008	$253.76	$1075.00	$750.71	$1545.40	0.33	$659.98	$1457.20	0.38
2009	$275.53	$1135.40	$596.45	$1422.40	0.45	$650.67	$1473.80	0.41
Total 2007–2009	$792.91	$3359.40	$2025.42	$4513.60		$1922.14	$4411.30	
Spending above full attention				$1154.20			$1051.90	
% Difference				25.6%			23.9%	

Notes: Results of simulations allowing premiums to adjust to changes in consumer behavior. Predicted OOP costs are cross-enrollee averages per enrollee per year, including premiums. "Full Attention" assumes all consumers make a new choice every year. "Inattention (Using Observed Shocks)" uses observed shocks to infer percent of attentive consumers and implied premium coefficient. "Inattention (Iterative Algorithm)" iterates to a point where predicted premium increases consistent with percent attentive consumers.

TABLE 9 Simulated per-Person Spending, Robustness

	P(attn.) = 0.2		P(attn.) = 0.4		P(attn.) = 0.6		P(att'n) = 0.8	
	Premium	OOP	Premium	OOP	Premium	OOP	Premium	OOP
2007	$1128.50	$1991.00	$639.37	$1507.60	$445.11	$1318.30	$337.10	$1215.90
2008	$1159.30	$1947.60	$633.46	$1431.50	$430.88	$1237.80	$320.36	$1134.90
2009	$1153.70	$1964.50	$661.66	$1484.30	$464.54	$1299.60	$351.14	$1198.90
Total 2007–2009		$5903.10		$4423.40		$3890.90		$3549.70
Spending above full attention		$2543.70		$1064.00		$496.30		$190.30
% Difference		43.1%		24.1%		12.9%		5.4%

Notes: Predicted OOP costs are cross-enrollee averages per enrollee per year, including premiums. Premiums include both basic and enhanced premiums. Each column assumes a different percent of consumers who are attentive, and hence a different premium coefficient in utility equation.

TABLE 10 Government Savings from Repricing Due to Full Attention

Year	NAMBA: Full Attn.	NAMBA: Inattention	γ_t	Annual ($) Savings/Enr.	Non-LIS Enrollment	Savings ($ million)
2006	$1027.40	$1027.40	0.65			
2007	$964.30	$1061.50	0.66	$64.15	8,120,524	$521 million
2008	$971.90	$1049.80	0.65	$50.88	8,413,202	$428 million
2009	$1001.20	$1071.50	0.64	$44.99	8,572,910	$386 million
Total 2007–2009						$1335 million

Notes: Results of program cost savings calculation. Columns 1 and 2 are predicted unweighted national average bids, PDP and MA plans nationally, measured in $ per year. We use the iterative algorithm corresponding to column 3 of Table 8 to simulate inattention. Per-enrollee average savings are the difference between the two average bids scaled by the proportion paid by the government. Non-LIS enrollment reported in national plan data generously provided by Francesco Decarolis. γ_t is defined in Section 8.

The third column of Table 8, labelled "Inattention (Iterative Algorithm)," contains our preferred specification under inattention. As described above, we begin with the outcome under "Inattention (Using Observed Shocks)" and iterate, predicting the shocks in iteration k using premiums generated in iteration $k-1$, and repeating until convergence. Thus, the predicted premiums reported in this column are consistent with the proportion of consumers assumed to be attentive (i.e., the proportion with a nonzero premium coefficient). The results change only slightly: the probability of attention is slightly higher than in column 2, varying from 0.38–0.42 in the years 2007–2009, and the corresponding consumer spending is slightly lower. Savings from a move to full attention are now $1051.90, or 23.9% of total spending.

Table 9 sets out additional robustness tests in the spirit of providing a range of estimates of the effect of consumer inattention. We report consumer spending under different assumed proportions of attentive and inattentive consumers in the population: we allow the proportion attentive to vary from 0.2 to 0.8, in increments of 0.2. The estimates are intuitive. The higher the proportion attentive, the higher the effective premium elasticity, and the lower is consumer spending on premiums, both because plans adjust their premiums strategically and because consumers increasingly prefer low-premium plans. The relationship between percent attentive and total consumer spending is not linear: the spending reductions fall with each incremental increase in percent attentive. This is unsurprising, because the premium coefficient is one input into a nonlinear choice model that determines plan market shares and hence, both bids and consumer spending. However, it is clear that the higher the initial percent attentive, the lower the saving from removing inattention. The estimates range from a 43% saving if only 20% of consumers are initially attentive, to a 5% saving if the initial proportion is 80%.

Table 10 reports the cross-plan unweighted average bids (the NAMBA) that generate the premiums reported in Table 8, with and without inattention. Because we assume that all consumers are attentive in 2006, the same average bid—$1027.40 per enrollee per year—is predicted in both columns 1 and 2 of the table ("attentive" and "inattentive"). In later years, as expected, average bids are lower when all consumers are assumed to be attentive (column 1) than under our preferred iterative algorithm for inattention (column 2). For example, in 2007, the predicted NAMBA under full attention is $964.30 per enrollee per year, whereas under inattention, the corresponding figure is $97 higher at $1061.50. The difference varies over years, with changes in plan costs and the observed enhanced premium, but it is greater than 6% of the "inattention" level in every year 2007–2009.[34] Finally, we note that the unweighted national average NAMBA observed in the data for 2006 is $1108; it varies between $965 and $1110 in the years 2007–2009. Although plan dynamic choices are likely to generate changes over time that we do not expect our simulations to match, the fact that the difference between observed and predicted NAMBA is never more than 9% of our simulated estimate is evidence of the accuracy of our cost measures.[35]

If we are willing to assume that our New Jersey estimates can be extrapolated to the entire nation, we can use the predicted bids to calculate the implied government savings for enrolled consumers over the years 2007–2009. These savings are substantial. Program cost savings result from the reduction in plan bids, of which the government pays a sizeable proportion. Applying the conservative assumption that reinsurance costs remain fixed so that the government saves a fraction of the difference in average bids equal to one minus the Base Beneficiary Percentage (γ_t in equation (5)), we find that government savings amount to between $45 and $64 per covered life per year in the years 2007–2009. Assuming further that low-income subsidy payments are unaffected and multiplying this figure by the non-LIS PDP population in each year generates an estimate of the government's total savings from reduced bids. We predict savings of $1.3 billion, or 1% of the government's cost of this part of the program over the three years 2007–2009.[36] This

[34] The reductions in bids reported here are smaller than the reductions in consumer spending due to endogenous consumer choice. The Table 10 bid numbers are unweighted averages across plans, consistent with the method used to calculate government program costs in the first few years of the program. In contrast, Table 8 reports averages across enrollees rather than plans. Consumers tend to choose lower-premium plans, particularly in the "full attention" simulations when the premium coefficient is most negative.

[35] These reported observed NAMBAs are unweighted national averages, whereas those in Table 7 are enrollment-weighted averages for New Jersey PDPs only (reported in $ per month for comparison with our cost estimates).

[36] This simple calculation assumes that, if inattention is removed nationally, the reduction in bids for PDP plans outside New Jersey will be the same as the average predicted reduction in New Jersey PDP plan bids, whereas MA plan bids remain constant. The $1.3 billion saving is 1% of the government's cost of subsidizing PDP premiums for non-LIS enrollees nationally.

is a conservative estimate: we use unweighted average bids, consistent with the use of unweighted averages to calculate the NAMBA in the first few years of the program, so our projections do not account for the impact on the weighted NAMBA of differential enrollee sorting into low-cost plans when bids change.[37] Clearly, we have made multiple assumptions to arrive at these numbers. However, when combined with the theoretical results discussed in Section 4, our estimates are sufficient to provide clear evidence that removing inattention would lead to substantial savings both to the federal government and to consumers.

The CBO calculated that from 2007–2010, the cost of the drug component of the basic benefit increased by 2.8% per annum on average for Part D enrollees while administrative costs and profits rose at 6.7% (CBO report, "Competition and the Cost of Medicare's Prescription Drug Program 2014"). Premiums increased from a weighted average of \$25.93 to \$37.25 from 2006 to 2010 according to the Kaiser Family Foundation (Hoadley, 2015). This is the environment we explore with our data and seek to explain as rational pricing in the face of consumer inattention. Interestingly, the environment changed significantly in the second five years of the program. From 2010 to 2014, the rate of generic penetration increased significantly and the introduction of new branded blockbuster drugs slowed. National pharmaceutical expenditure actually fell in nominal terms in 2012 and 2013, according to IMS data (Schumock et al., 2014). However, from 2010 to 2015, stand-alone Part D premiums stayed approximately constant at between \$37.02 and \$38.54 (Hoadley, 2015). Because drug costs fell modestly in those years, it is not clear what happened to the margins of the insurers participating in Part D. Ideally, effective regulation should ensure that costs are reflected in prices. That is, falling drug costs should benefit consumers in the form of falling premiums if a program is delivering the competitive benefits society expects. The available evidence suggests that this did not happen, adding weight to our hypothesis that consumer inattention limits the effectiveness of competition in this market. It may be the case that the underlying cost environment has changed again recently, as several sources (e.g., the Express Scripts 2014 Drugs Trend Report) report significant increases in specialty drug spending for 2014. Further research into the impact of consumer choice in Part D on competition is clearly warranted.

9. Conclusions

■ In this article, we have developed a model of consumer choice in the Part D program and have analyzed how firms set prices in response to the presence or absence of consumer inattention. The data support a model where consumers face costs of processing information. This leads them to avoid making new choices, rolling over their plan selections from one year to the next unless shocked by a change to their current plan or their current health.

We provide evidence that firms' premium choices are responsive to consumers' search frictions. In particular, when consumers are attentive, firms are incentivized to lower their margins, resulting in lower premiums. Using our estimates of consumer behavior and a model of firm price-setting, we use a simple approach to approximate the steady state effects of consumer inattention on premiums charged and consumer spending. We predict a large price response to full consumer attention because, in that environment, an effective way for plans to attract customers is by lowering premiums. Our simulations indicate that the combination of the demand- and supply-side changes would reduce the amount consumers spend by 24%. The natural plan response of increasing other components of the price, like the OOP cost schedule, is constrained by the quite tightly regulated standard benefit levels.

We do not consider price dynamics. We note that theory predicts an increasing path of prices, particularly in early years of the program, as plans respond to incentives to "invest" in new

[37] This is the primary reason why the predicted government saving is lower than the weighted average saving to consumers from removing inattention. A second reason for the difference in *percent* savings is that the denominator for the government is the total cost of subsidizing PDP premiums, including risk payments and reinsurance, whereas for the consumer it is simply out-of-pocket payments.

enrollees and then "harvest" their installed base; this implies that the price effects of inattention may be increasing over time. However, there are also other, potentially offsetting effects, such as insurer learning about the Part D environment during the first few years of the program. Steady state equilibria could also involve plans cycling between high and low prices, particularly in situations where several plans are offered by a single insurer. Our simulations do not capture these effects, nor do they consider the impact of changes in the number of plans per insurer, or the number of insurance carriers, active in each market. All of these issues are potentially fruitful areas for future research. However, our approach is sufficient to provide a feasible range of estimates of the price effects of inattention that demonstrates the empirical and policy relevance of the issue.

The role of plan repricing in response to more frequent and effective consumer search has not been analyzed to the best of our knowledge in the Medicare Part D economics literature to date. It is an important element in the evaluation of any policy that would help consumers choose better plans. In particular, although clearly the extrapolation of our New Jersey estimates to the national level should be interpreted with some caution, the implied government savings from consumer choice—$1.3 billion in the years 2007–2009—indicate how important well-designed insurance marketplaces can be. Indeed, without effective consumer choice that puts market pressure on insurers, a policy of privatizing the delivery of benefits can be very expensive. This cost of privatization should be taken into account by policy makers in light of our results.

References

Abaluck, J. and Gruber, J. "Heterogeneity in Choice Inconsistencies Among the Elderly: Evidence from Prescription Drug Plan Choice." *American Economic Review*, Vol. 101 (2011), pp. 377–381.

———. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." NBER Working Paper no. 19163, 2013.

Agarwal, S., Driscoll, J., Gabaix, X., and Laibson, D. "The Age of Reason: Financial Decisions over the Life Cycle and Implications for Regulation." *Brookings Papers on Economic Activity*, (2010), pp. 51–101.

Beggs, A., and Klemperer, P. "Multiperiod Competition with Switching Costs." *Econometrica*, Vol. 60 (1992), pp. 651–666.

Berry, S. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, Vol. 25 (1994), pp. 242–262.

Busse, M., Pope, D., Pope, J., and Silva-Risso, J. "The Psychological Effect of Weather on Car Purchases." *The Quarterly Journal of Economics*, Vol. 130 (2014), pp. 371–414.

Busse, M., Simester, D., and Zettelmeyer, F. "The Best Price You'll Ever Get: The 2005 Employee Discount Pricing Promotions in the US Automobile Industry." *Marketing Science*, Vol. 29 (2010), pp. 268–290.

Cabral, L. and Fishman, A. "Business as Usual: A Consumer Search Theory of Sticky Prices and Asymmetric Price Adjustment." *International Journal of Industrial Organization*, Vol. 30 (2012), pp. 371–376.

Carey, C. "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D." *American Economic Journal: Economic Policy*, (2017).

Cebul, R., Rebitzer, J., Taylor, L., and Votruba, M. "Unhealthy Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance." *American Economic Review*, Vol. 101 (2011), pp. 1842–1871.

Dafny, L., Ho, K., and Varela, M. "Let Them Have Choice: Gains from Shifting Away from Employer-Sponsored Health Insurance and Toward an Individual Exchange." *American Economic Journal: Economic Policy*, Vol. 5 (2013), pp. 32–58.

Decarolis, F. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review*, Vol. 105 (2015), pp. 1–34.

Decarolis, F., Polyakova, M., and Ryan, S. "The Welfare Effects of Supply-Side Regulations in Medicare Part D." NBER Working Paper no. 21298, 2016.

Dube, J., Hitsch, G., and Rossi, P. "State Dependence and Alternative Explanations for Consumer Inertia." *RAND Journal of Economics*, Vol. 41 (2010), pp. 417–445.

Duggan, M. and Scott Morton, F. "The Effect of the Medicare Drug Benefit on Pharmaceutical Prices and Utilization." *American Economic Review*, Vol. 100 (2010), pp. 590–607.

Einav, L., Finkelstein, A., and Polyakova, M. "Private Provision of Social Insurance: Drug-Specific Price Elasticities and Cost Sharing in Medicare Part D." NBER Working Paper no. 22277, 2016.

Ericson, K. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy*, Vol. 6 (2012), pp. 38–64.

———. "When Consumers Do Not Make An Active Decision: Dynamic Default Rules and Their Equilibrium Effects." NBER Working Paper no. 20127, 2014.

FARRELL, J. AND KLEMPERER, P. "Coordination and Lock-In: Competition with Switching Costs and Network Effects." In *The Handbook of Industrial Organization*. New York: North Holland, 2007.

GREENWALD, L. AND WEST, N. "Medicare Part D Demonstration Focus Group Report." RTI Project no. 02007964.023.006, 2007.

HANDEL, B. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, Vol. 103 (2013), pp. 2643–2682.

HECKMAN, J. "Identifying the Hand of the Past: Distinguishing State Dependence from Heterogeneity." *American Economic Review*, Vol. 81 (1991), pp. 75–79.

HEISS, F., LEIVE, A., MCFADDEN, D., AND WINTER, J. "Plan Selection in Medicare Part D: Evidence from Administrative Data." *Journal of Health Economics*, Vol. 32 (2013), pp. 1325–1344.

HEISS, F., MCFADDEN, D., AND WINTER, J. "Who Failed to Enroll in Medicare Part D, and Why?" *Health Affairs*, Vol. 25 (2006), pp. 344–354.

HOADLEY, J. "Medicare Part D at 10 Years." Kaiser Family Foundation Issue Brief, 2015.

HOADLEY, J., HARGRAVE, E., SUMMER, L., CUBANSKI, J., AND NEUMAN, T. "To Switch or Not to Switch: Are Medicare Beneficiaries Switching Drug Plans to Save Money?" Kaiser Family Foundation Issue Brief, 2013.

HOADLEY, J., SUMMER, L., AND HARGRAVE, E. "Medicare Part D in Its Ninth Year: The 2014 Marketplace and Key Trends, 2006–2014." Kaiser Family Foundation Issue Brief, 2014.

HONKA, E. "Quantifying Search and Switching Costs in the US Auto Insurance Industry." *RAND Journal of Economics*, Vol. 45 (2014), pp. 847–884.

HORTACSU, A., MADANIZADEH, S., AND PULLER, S. "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market." NBER Working Paper no. 20988, 2015.

ILLANES, G. "Switching Costs in Pension Plan Choice." Northwestern University Working Paper, 2016.

IYENGAR, S. AND KAMENICA, E. "Choice Proliferation, Simplicity Seeking, and Asset Allocation." *Journal of Public Economics*, Vol. 94 (2010), pp. 530–539.

KETCHAM, J., KUMINOFF, N., AND POWERS, C. "Estimating the Heterogeneous Welfare Effects of Choice Architecture: An Application to the Medicare Prescription Drug Insurance Market." NBER Working Paper no. 22732, 2016.

KETCHAM, J., LUCARELLI, C., MIRAVETE, E., AND ROEBUCK, M. "Sinking, Swimming, or Learning to Swim in Medicare Part D." *American Economic Review*, Vol. 102 (2012), pp. 2639–2673.

KETCHAM, J., LUCARELLI, C., AND POWERS, C. "Paying Attention or Paying Too Much in Medicare Part D." *American Economic Review*, Vol. 105 (2015), pp. 204–233.

KLEMPERER, P. "Markets with Consumer Switching Costs." *The Quarterly Journal of Economics*, Vol. 102 (1987), pp. 375–394.

KLEMPERER, P. "Competition when Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics and International Trade." *Review of Economic Studies*, Vol. 62 (1995), pp. 515–539.

KLING, J., MULLAINATHAN, S., SHAFIR, E., VERMEULEN, L., AND WROBEL, M. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *The Quarterly Journal of Economics*, Vol. 127 (2012), pp. 199–235.

LAVETTI, K. AND SIMON, K. "Strategic Formulary Design in Medicare Part D Plans." NBER Working Paper no. 22338, 2016.

LEVINSON, D. "Higher Rebates for Brand-Name Drugs Result in Lower Costs for Medicaid Compared to Medicare Part D." Department of Health and Human Services, 2011.

LUCO, F. "Switching Costs and Competition in Retirement Investment." Working Paper, 2016.

MILLER, K. "Do Private Medicare Firms Have Lower Costs?" Working Paper, 2014.

POLYAKOVA, M. "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D." *Applied Economics*, Vol. 8 (2016), pp. 165–195.

RADNER, A. "Viscous Demand." *Journal of Economic Theory*, Vol. 112 (2003), pp. 189–231.

REIS, R. "Inattentive Consumers." *Journal of Monetary Economics*, Vol. 53 (2006), pp. 1761–1800.

SALLEE, J. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics*, Vol. 57 (2014), pp. 781–820.

SCHUMOCK, G., LI, E., SUDA, K., MATUSIAK, L., HUNKLER, R., VERMEULEN, L., AND HOFFMAN, J. "National Trends in Prescription Drug Expenditures and Projections for 2014." *American Journal of Health-System Pharmacy*, Vol. 71 (2014), pp. 485–499.

SIMS, C. "Implications of Rational Inattention." *Journal of Monetary Economics*, Vol. 50 (2003), pp. 665–690.

STOCKING, A., BAUMGARDNER, J., BUNTIN, M., AND COOK, A. "Examining the Number of Competitors and the Cost of Medicare Part D." CBO Working Paper Series 2014-04, 2014.

SULLIVAN, K. "How to Think Clearly About Medicare Administrative Costs: Data Sources and Measurement." *Journal of Health Politics, Policy and Law*, Vol. 38 (2013), pp. 479–504.

SUMMER, L., HOADLEY, J., AND HARGRAVE, E. "The Medicare Part D Low-Income Subsidy Program: Experience to Date and Policy Issues for Consideration." Kaiser Family Foundation Policy Brief, 2010.

TRAIN, K. *Discrete Choice Methods with Simulation*. New York: Cambridge University Press, 2009.

VARIAN, H. "A Model of Sales." *American Economic Review*, Vol. 70 (1980), pp. 651–659.

WEINTRAUB, G., BENKARD, C., AND VAN ROY, B. "Markov Perfect Industry Dynamics with Many Firms." *Econometrica*, Vol. 76 (2008), pp. 1375–1411.

Supporting information

Additional supporting information may be found in the online version of this article at the publisher's website:

Appendix Table 1: Defined Standard Benefit Parameters, 2006–2013
Appendix Table 2: Sample Composition
Appendix Table 2B: Age Distribution
Appendix Table 2C: Part D Tenure
Appendix Table 3: Average Plan Quality
Appendix Table 4: Switching by Demographic Group
Appendix Table 5: Spending Gap by Switch Decision
Appendix Table 6: Proportion Within X% of Lowest-Cost Plan
Appendix Table 7: Probit Regressions on Switch Decision
Appendix Table 8: Next-Year Plan Choices and Overspending by Shock, Switchers Only
Appendix Table 9: Following-Year Plan Characteristics Choices, Switchers and Non Switchers
Appendix Table 10: Predicted Overspending Regressions
Appendix Table 11: Estimated Structural Demand Coefficients
Appendix Table 12: Descriptive Statistics for Demand Model Variables
Appendix Table 13: The Five Largest New Jersey PDP Plans, 2006
Appendix Table 14: Premium Dispersion in New Jersey DSB Plans
Appendix Table 15: Annual Premium Increases, Accounting for "Sister" Plan Entry