NBER WORKING PAPER SERIES

THE WELFARE EFFECTS OF SUPPLY-SIDE REGULATIONS IN MEDICARE PART D

Francesco Decarolis Maria Polyakova Stephen P. Ryan

Working Paper 21298 http://www.nber.org/papers/w21298

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 June 2015

We are grateful to Mark Duggan, Liran Einav, Amy Finkelstein, Jon Gruber, Ben Handel, Kate Ho, Thomas Jeitschko, Claudio Lucarelli, Aviv Nevo, Ariel Pakes, Nancy Rose, Marc Rysman, and numerous seminar participants. Decarolis is grateful to the Sloan Foundation (grant 2011-5-23 ECON) for financial support. We also gratefully acknowledge support from the National Science Foundation (SES-1357705). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peerreviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Francesco Decarolis, Maria Polyakova, and Stephen P. Ryan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Welfare Effects of Supply-Side Regulations in Medicare Part D Francesco Decarolis, Maria Polyakova, and Stephen P. Ryan NBER Working Paper No. 21298 June 2015, Revised November 2016 JEL No. H2,H4,I11,I18,L1,L2

ABSTRACT

The efficiency of publicly-subsidized, privately-provisioned social insurance programs depends on the interaction between insurer behavior and public subsidies. We study this interaction within Medicare Part D Prescription Drug Plan (PDP) markets. Using a structural model of supply and demand, we find: consumers purchase too few and too socially-costly PDP plans; insurers price near marginal cost; the primary driver of welfare is the opportunity cost of government spending on other Medicare programs; and the current subsidization policy achieves a level of total welfare close to that obtained under an optimal in-kind subsidy, but is far from the social planner's first-best solution.

Francesco Decarolis Department of Economics Boston University 270 Bay State Rd. Boston, MA 02215 fdc@bu.edu

Maria Polyakova Stanford School of Medicine Redwood Building T111 150 Governor's Lane Stanford, CA 94305 and NBER maria.polyakova@stanford.edu Stephen P. Ryan Department of Economics University of Texas at Austin BRB 3.134D 2225 Speedway Stop C3100 Austin, TX 78712-1690 and NBER sryan@utexas.edu

1 Introduction

Subsidizing the private provision of health insurance constitutes a large and growing fraction of government's expenditures in health care. Recently, the design of complex subsidization policies within the Affordable Care Act has spurred extensive political debate and become a focal point of disagreement around the Act. Potentially even more important, however, is the ongoing shift from the traditional fee-for-service model to subsidized private provision of Medicare coverage, as the budgetary outlays on the Medicare program are roughly five times the projected annual spending on the ACA Exchanges, and amounted to more than \$500 billion in 2014.¹ This development substantially changes how the government spends money in health care, pivoting away from direct reimbursement of physician and hospital services toward the subsidization of private insurance plans. Despite the importance of subsidy spending in budgetary outlays, very little is known about the efficiency and the distributional effects of the existing subsidization policies.

In this paper we use the institutional environment of Medicare Part D, a privately provided, publicly-subsidized insurance program for prescription drugs, to derive lessons about the efficiency implications of different subsidy designs. Part D is an important, controversial, and expensive program, with federal spending totaling more than \$76 billion annually. We focus our study on the mechanism determining subsidies for Prescription Drug Plans (PDP) in Medicare Part D, particularly on strategic behavior by insurers. Despite the importance of supply-side incentives in this setting, the academic and policy debate has so far mostly focused on individuals' choices in the Part D program, leaving aside insurers' pricing incentives. To start closing this gap, we explore the allocative efficiency implications of the mechanism that Medicare uses to determine subsidies in Medicare Part D. This mechanism relies heavily on insurers' behavior, deriving government subsidies from prices set by insurers. Since subsidies account for the majority of insurers' Part D revenues, the endogeneity of subsidies to insurers' pricing raises concerns about potential efficiency distortions on the supply-side.

Our research strategy starts with the estimation of demand for prescription drug plans. In each market, firms offer a list of insurance plans which vary across several dimensions such as the size of the deductible, the set of drugs that are covered, whether the plan has a "donut-hole," which is a region of expenditures for which the plan reverts to 100 percent

¹Source: Congressional Budget Office, 2014 Medicare Baseline. In terms of federal spending, out of total federal outlays of \$3.5 trillion in 2013, the net federal outlays for Medicare amount to 14 percent or \$492 billion.

co-insurance, and the plan's premium. Demand in Part D is slightly more complicated than the typical setting due to the presence of two groups of consumers: so-called regular enrollees and low-income (LIS) enrollees.² Regular enrollees make unrestricted choices from all plans offered in their region and pay a partially-subsidized premium. In contrast, low-income enrollees, who constitute 35 percent of all enrollees, are randomly assigned to eligible plans by the Centers for Medicare and Medicaid Services (CMS) and pay nothing. These enrollees, however, can and do opt out of the random assignment process and freely choose any plan at additional cost.³ Using four years of data on the characteristics and enrollments of all Part D PDP plans across all 34 Medicare Part D regions in the US, we estimate demand for both regular and LIS enrollees using the random coefficients discrete choice framework pioneered by Berry, Levinsohn, and Pakes (1995).

Given demand estimates for insurance plans, we then turn our attention to modeling the behavior of insurers. A critical piece of this puzzle is the rule for how a firm's pricing decision, hereafter referred to as its bid, are turned into premiums that enrollees face. Medicare beneficiaries do not face full prices or bids set by insurers; instead, there is an intermediate process by which CMS decides on how much of the insurer's bid will be paid by the government in subsidies, and how much will be paid by enrollees in premiums. In this process, CMS takes the sum of all bids for all participating insurers in the US, combines them with bids from the Medicare Advantage market, averages them using enrollment weights from the previous year, and takes a fraction of the resulting number to obtain the base subsidy. The premium of a given plan is then determined by taking the maximum of zero and the firm's bid minus this base subsidy. This pricing mechanism has three effects on market outcomes. First, consumers face premiums that are strictly lower than firm bids, which increases demand. Second, the relative premiums of plans are distorted by this mechanism; this is important since it distorts the choices behavior of consumers across plans. Third, the same bids determine the plans' eligibility to enroll the randomly-assigned LIS enrollees. Only plans with a premium below the average premium in their market are eligible for random assignment of LIS enrollees. Consequently, there is a key linkage between the two groups: the bidding process by which plans qualify to be eligible for low-income assignments also influences premiums for regular enrollees. Thus, these incentives distort both the public payments for low-income enrollees and the prices and choices of regular enrollees. We model firm behavior in light of these incentives and recover estimates of plan-level marginal costs.

²Center for Medicare and Medicaid Services (2009) explains how LIS eligibility is determined. In general, consumers below some multiple of the federal poverty level are eligible for LIS classification.

³For example, in 2011 about one-third of LIS enrollees had opted out of the random assignment system.

With demand and supply cost estimates in hand, we then characterize the welfare effects of the current subsidy mechanism. Our welfare estimates depend on the estimated consumer surplus, producer profits, and the social cost of government spending. We assume that the deadweight loss of taxation is given by 30 cents per dollar of revenue raised (Hausman and Poterba, 1987). We also make two critical assumptions in computing welfare. First, we initially assume that the rest of the world does not change as we modify the subsidy mechanism in Part D PDP.⁴ As such, all of our counterfactual results are subject to the usual partialequilibrium critiques. Second, all of our estimates, demand, marginal cost, and government spending, are measured relative to their opportunity cost. Consumers in this market are not left without coverage options if the Part D PDP market were to shut down; one can readily see this as the inside share of consumers in Part D PDP is only 38 percent in 2010, which is the last year of our data. The remaining 62 percent are primarily covered by private insurance or a similar insurance program offered under Medicare Advantage (MA-PD). Evidence from consumer level data indicates that out of the few enrollees that switch out of the PDP plans, two-thirds move to an MA-PD plan. Producers face a direct marginal cost of providing the good here, but also the opportunity cost of potentially serving the same consumer in the MA-PD market. Indeed, about 90 percent of PDP plans are offered by insurers that also offer an MA-PD plan. The government spending opportunity cost is particularly salient, as we conservatively assume that consumers would substitute from Part D PDP plans exclusively to MA-PD plans rather than dropping a publicly subsidized program altogether. This implies that all of our estimates—demand, marginal cost, government spending, and, thus, social welfare—are relative to the outside option.

We first calculate welfare estimates for the observed prices and allocations. In the market for regular enrollees, our findings suggest that relative to the existing outside option, the current levels of subsidies in the stand-alone Prescription Drug Plans are generating negative nominal welfare with a return of only 17 cents of surplus for every dollar of government spending. However, once the foregone costs of providing similar services in MA-PD are considered, the program generates substantial surplus, with a return of \$1.11 per dollar of opportunity cost. This is one of our primary findings: the positive welfare effect of Part D PDP is driven almost exclusively by the opportunity cost of government spending. On its own merits, the total cost of providing subsidized goods exceeds their benefits; expenditures of \$9.4 billion generated \$2.5 billion of consumer surplus and \$450 million of producer profit. However, we estimate that foregone costs of providing similar coverage in MA-PD (under

⁴We relax this assumption in a separate set of counterfactuals below.

partial equilibrium assumptions) is \$8.3 billion. Considering the opportunity cost and the deadweight loss of taxation to raise government funds, we estimate that the program in its current form generates \$1.5 billion in surplus. In the LIS enrollee market, similar economic forces are at play, with an even larger opportunity cost of government spending that increases surplus to \$3.1 billion.

To understand the role of the various pieces of the subsidy mechanism, and to assess its efficiency relative to several alternative approaches, we perform several counterfactual experiments. We start with two counterfactuals that examine regulatory links between parts of the Part D subsidy mechanism. Recognizing potential problems arising from mixing together the regular enrollees and the LIS enrollees, several policy initiatives have proposed removing the LIS enrollees to their own market. We simulate this policy and we find that consumer surplus and producer profit increase relative to the observed mechanism, but overall surplus declines slightly, as the net surplus generated by the marginal consumers is exceeded by the social cost of subsidizing the program. This finding is further exacerbated if we were to remove MA-PD plans from the weighted average bid used to compute the baseline PDP subsidy. The message from these two counterfactuals is that these links help reduce the generosity of the subsidy, which in turn leads to improvements in welfare. We take the latter counterfactual as our baseline for further investigation of the properties of the mechanism, as it removes strategic incentives for competing for LIS enrollees and possible feedback effects in the MA-PD market.

One of the primary motivations for managed competition is that firms will compete for consumers, driving down prices. To assess the competitiveness of the market, we perform two counterfactuals where we change the ownership structure. In the first, we assume that each plan is its own firm; in the second, we assume that every plan in each market belongs to one firm. Compared to the baseline counterfactual, we find the expected pattern that profits increase greatly and consumer surplus declines under the monopolistic regime, with the opposite pattern under atomistic competition. Interestingly, total surplus declines under both scenarios, but decreases slightly more under atomistic competition. Although the monopolist increases prices, this generates welfare increases as the marginal consumers who exit the market had valuations for PDP plans that were below their social marginal cost.

This highlights a general tension in this setting: the social planner must balance the benefits of additional consumer surplus and producer profits against the social cost of subsidizing the provision of those goods. To formalize this, we perform several counterfactuals where the government sets prices and/or plan characteristics directly. In the first, prices are set at private marginal cost. In the second, the government acts as the social planner, maximizing total welfare. In the third, the government sets all plan characteristics equal to the lowest cost plan in each market and sets premiums at marginal cost net of observed subsidy level.

Under marginal cost pricing, consumer surplus is less than half of the current mechanism, driven by a more than doubling of consumer premiums and a corresponding precipitous decline in the amount of consumers choosing to buy a Part D PDP plan. This is not a completely unexpected result; on the one hand, prescription drug coverage in general is certainly a valuable product for seniors. For example, Town and Liu (2003) conclude in their estimates of welfare effects from the introduction of Medicare Advantage program that the prescription drug insurance part of the program was extremely valuable for the Medicare population. At the same time, Engelhardt and Gruber (2011) find evidence of substantial crowd-out, where Part D insurance was used merely as a substitute for other prescription drug coverage sources. Given the outside option, we may have expected to see a large substitution to the outside good if consumers faced the marginal cost.

Setting prices at marginal cost, which is usually the benchmark for welfare maximization, produces suboptimal outcomes by ignoring an important component of welfare: the opportunity cost of government spending. To incorporate that dimension, we compute the social planner's problem. As expected, the social planner has high total surplus of \$3.6 billion, or almost three times the current mechanism. Enrollment in Part D PDP under the social planner is nearly 50 percent of the market. Consumer surplus is slightly higher than the observed mechanism, but the distribution of equilibrium prices is completely different. The social planner highlights and solves two inefficiencies of the current mechanism: average prices are lowered to induce more people to consume in PDP, but the relative prices are also changed to induce more consumers to choose less-costly plans. The counterfactual where the government sets all characteristics of plans equal to the plan with the lowest marginal cost in each market generates a large positive amount of welfare, \$2.8 billion. This counterfactual is able to attract a higher number of consumers at a relatively low cost into the PDP market. This counterfactual is also attractive for its simplicity, conditional on the government's ability to replicate or procure a low cost plan.

With these benchmarks in mind, we then proceed to investigate a menu of counterfactual subsidy-setting policies that CMS could implement in lieu of the current bid averaging process. The simplest scenario would be to provide fixed vouchers that could be used to buy a plan in the Part D market. We find that the current system operates like a voucher, in that the average bid mechanism is set by bids of all plans, and any individual firm has little influence on that average. Unsurprisingly, we can replicate the observed surplus very closely using a fixed voucher. Interestingly, the current system gives welfare very close to the optimal uniform voucher. Bridging the gap between a uniform voucher at the national level and the social planner's plan-specific prices, we also evaluate the welfare gains of instituting vouchers that vary at the regional level, but find that the welfare increase is very minor.

A second option would be to use a uniform proportional discount on all plans' bids. We find that proportional subsidies are, in general, a disastrous idea as firms simply scale their bids in proportion to the subsidy. Consumers face increasingly low premiums, firms are paid increasingly large bids, and government expenditures explode. That combination results in large negative welfare losses.

These counterfactuals highlight the fundamental economic forces at work in the Part D PDP market. One important assumption that we make in these counterfactuals is that the value of the MA-PD outside option does not change with changes to the PDP market. While this is a useful heuristic device for isolating the interactions of firms, consumer preferences, and the subsidy mechanism in order to increase our understanding of how the market works, it is also reasonable to expect that CMS would implement analogous changes to the MA-PD market at the same time as it changed the PDP market. We address this by allowing the average value of the outside option to change with the change in the average generosity of the subsidy in the PDP market. The takeaway from this exercise is that outcomes are similar to the baseline case but are generally exaggerated by the endogenous adjustment to the MA-PD market. The intuition is that making the subsidy in both parts of the Part D program more generous increases the overall cost of providing coverage. Since the value of coverage has not changed, this results in a decline in welfare. This can be seen most clearly in the case of vouchers, where this logic is taken to an extreme and the optimal voucher is \$0.

Our baseline results remove the LIS enrollees from the market when computing our counterfactuals. To assess the influence of adding these consumers back into the market, we evaluate a matrix of counterfactuals where we compute outcomes under different levels of vouchers given to regular and LIS enrollees. The optimal voucher does not change for the regular enrollees, and we find that the optimal level of vouchers for LIS enrollees is very low. The intuition is similar to other cases: increasing the voucher for LIS enrollees increases consumer surplus and enrollment in that group, but the social cost exceeds the value the marginal enrollee places on the services.

Our paper is related to a large theoretical literature that has examined the role and motivation for in-kind subsidies in different sectors of the economy; surprisingly, however, the empirical analysis of the motivation and effects of such government policies is much less explored (Currie and Gahvari, 2008). In health insurance, the literature has investigated the effects of tax subsidies to employer-provided health insurance, for example in Gruber and Washington (2005); moreover, in the classic illustration of an adverse selection spiral, Cutler and Reber (1998) discuss the role of subsidy design (by the employer) in employersponsored plans. The recent expansion of federal health insurance programs—from Medicare and Medicaid to the ACA health insurance exchanges—into private markets has shifted the large public policy interest to how the federal budget subsidizes these programs. For example, Enthoven (2011) and Frakt (2011) discuss some of the key conceptual points and the policy debate. Conceptually and methodologically, our paper is closest to Curto et al. (2015) that explores the questions about subsidies, competition, and market design in the context of Medicare Advantage.

The current paper is also related to the growing literature that analyzes the Medicare Part D program as a prominent example of a health insurance program with consumer choice. This literature has so far mostly focused on demand questions. Several papers have explored the rationality of individual choices (Heiss et al., 2010, 2013; Abaluck and Gruber, 2011, 2016; Ketcham et al., 2012; Kesternich et al., 2013; Kling et al., 2012; Vetter et al., 2013; Winter et al., 2006; Ketcham et al., 2015). Relatedly, Ericson (2014); Miller and Yeo (2014); Abaluck and Gruber (2016); Ho et al. (2015); Polyakova (2016) explore the presence and role of inertia in the individual choices of Part D contracts. Einav et al. (2015) study the effect of non-linear contract structure on the drug consumption decisions in Part D. Related work has explicitly considered dynamic incentives within Part D contracts (Abaluck et al., 2015; Gowrisankaran et al., 2015). A number of papers, in economics and health services research, have examined the effect of Part D on drug utilization, adherence, and health outcomes for the elderly, for example Ketcham and Simon (2008).

Further, this paper is related to a substantial theoretical and empirical literature on the supply-side effects of government regulation. Laffont and Tirole (1993) gives a classic reference on the multitude of theoretical issues. Our research question is related to the issues of government procurement in health care, such as Duggan (2004) and Duggan and Scott Morton (2006). Ericson (2014) raises the questions of insurer strategies in Part D, arguing that insurers are exploiting individual inertia in their pricing decisions. Duggan et al. (2008); Duggan and Scott Morton (2010) estimate the effect of Part D on drug prices, and Yin and Lakdawalla (2015) analyzes how Part D enrollment affects private insurance markets. Decarolis (2015) focuses entirely on the supply-side, documenting that insurers are pricing strategically to take advantage of low income subsidy policies in Part D.

There are a number of papers using data and methodology similar to the present paper, although they focus on very different research questions. Lucarelli et al. (2012) is primarily concerned with the number of plans that are offered in each market. They consider two different counterfactuals along those lines: capping the number of plans that each firm can offer in each market at two, and eliminating all plans that have coverage in the donut hole. Ho et al. (2015) focus on measuring the equilibrium effects of choice frictions in Part D. They document that consumers are inattentive, which lowers the residual demand elasticity that suppliers face when pricing. In a related recent working paper, Wu (2016) poses and estimates a dynamic model of strategic supplier response to switching costs. Echoing Dubé et al. (2009), she finds that the equilibrium effect reduces prices. She studies a counterfactual of capping premium increases to mitigate supplier incentives to raise prices on captured consumers. Miller (2016) examines the effect of the interaction of the LIS subsidy with Medicare Part D's risk adjustment mechanisms on equilibrium prices. His paper focuses on the role of three risk adjustment mechanisms (individual risk adjustments, risk corridors, and reinsurance) as explaining the increase in drug prices after a reform in 2011. Miller and Yeo (2015) considers the introduction of a "public option" into Medicare Part D, while Chorniy et al. (2014) examines the effects of mergers among suppliers. While each of these papers uses similar data and methodological approaches to the present paper, none of them examine the subsidy mechanism's role on equilibrium outcomes or consider counterfactuals related to it.

The remainder of the paper is organized as follows: Section 2 discusses the key economic concepts. Section 3 describes the institutional details of the Medicare Part D market and our sources of data. Section 4 introduces the theoretical model underpinning our analysis, while Section 5 describes our empirical application of that model to the data and our results. Section 6 discusses our counterfactual pricing mechanisms and presents our results. Section 7 concludes.

2 Conceptual Framework

In imperfectly-competitive markets with differentiated products, such as Medicare Part D, subsidy policies create incentives that affect both consumer and producer behavior. To

illustrate the basic intuition behind these incentives, we consider the following simple model of subsidized competition with differentiated products. Suppose there are two firms, each selling one product to a unit mass of consumers. The utility to consumer *i* from product *j* is given by $u_{ij} = \delta_j - \alpha p_j + \epsilon_{ij}$, where δ_j is a measure of desirability of product *j* and p_j is its price. The outside option gives zero utility. The idiosyncratic taste shock, ϵ_{ij} , is distributed type I extreme value with the usual scale normalization. Consumers purchase the good that gives them the highest utility. Firms set prices to maximize profits; under the assumption that marginal costs are zero, the first-order condition for the price of *j*-th product is $p_j = 1/(\alpha(1-s_j))$, where s_j is the share of product *j*: $s_j = \exp(\delta_j - \alpha p_j)/(1 + \sum_{k=1}^2 \exp(\delta_k - \alpha p_k))$.

Let us now explore how subsidization may affect this economy. First, consider the case where the government introduces a proportional subsidy that reduces the effective price by 1-z; that is, the price in the utility function is effectively multiplied by $\hat{\alpha} = (1-z) \cdot \alpha$. The government raises funds for the subsidy through distortionary taxes that create a deadweight loss, $\lambda - 1$, for each dollar of subsidy provided. For our illustrative examples below, we set $\alpha = 1, \lambda = 1.3$, and $\delta = \{1.2, 1\}$, so that the first product is more desirable to consumers; in an insurance context, this would be a plan with more generous coverage.

Figure 1a illustrates the effects of increasing z from zero to one. For consumers, government subsidies distort choices along two margins. First, subsidizing a market makes it more attractive to marginal consumers, who will purchase a product only when prices are subsidized; this is shown in the figure as a decline in the outside option as the subsidy becomes more generous. This distorts consumption choices on the extensive margin *across* markets, and has important implications for welfare, as it can lead to consumers with relatively low willingness to pay purchasing goods with social costs above their valuations. A second, more subtle, margin is that the subsidy may influence both the marginal and inframarginal consumers if it changes the relative demand for products within the market. In general, consumers substitute from the now relatively more expensive plans to those that are now relatively cheaper as the subsidy mechanism changes their relative prices. This is shown in the figure as the increase in relative inside shares, computed as the share of product one over the share of product two. Here, consumers substitute to the more generous plan in equilibrium when subsidies are higher. Hence, subsidies may distort the allocative efficiency *within* the market.

The subsidy also affects the supply side of the market. Firms know that consumers have less elastic demand curves, and will charge higher markups as a result. To measure this effect, we define "excess price" as the percentage of the proportional subsidy which is set as a higher price by firms. For example, if the proportional subsidy is 10 percent, and the effective prices that consumers face decline by only 8 percent, we define the excess price to be 100 * (0.10 - 0.08)/0.10, or 20 percent. The excess price is plotted as a percentage against the right y-axis. The values range from zero when the subsidy is zero and one, to a little over six percent when the subsidy is just under 60 percent. This measures subsidy incidence and is the supply-side response to a demand-side subsidy analog of pass-through when a firm faces higher costs.

The form of the subsidy also matters. Figure 1b shows the same quantities with a prespecified flat "voucher" subsidy instead of a proportional subsidy. Consumers are given a voucher ranging from zero to one and a half to spend on an inside option. Increasing the generosity of the voucher decreases the share of the outside option as the equilibrium price that consumers face falls. While the ratio of the shares of the two inside options changes, it does so in the opposite direction from the proportional subsidy; the market share of the first (more valued) good falls as the voucher becomes more generous. We redefine the excess price in this case to be the percentage of the voucher that is not expressed in the decline of the consumer-facing price. For example, if the voucher is equal to 1, but the price only drops by 0.7, the excess price is 30 percent. In our simple model, firms are able to raise prices, so that the excess price is approximately 20 percent in response to the voucher is non-monotonic.

This discussion highlights that the government faces a complex, nonlinear set of economic forces when considering both how generously to subsidize markets and also how to provision those subsidies. Figure 2a shows the constituent components of the social planner's welfare maximization problem—consumer surplus, producer profits, government expenditures—along with total surplus, for the case of proportional subsidies. Figure 2b repeats the exercise for the flat voucher subsidy. In both cases, the optimal subsidy occurs in the interior, where the social planner trades off increases in consumer surplus and producer profit against socially-costly government expenditures. The figures also highlight one economic rationale for subsidization: market power. In our model, oligopolists set prices above marginal cost. Consistent with the theory of the second best, subsidizing demand leads to increased consumption in equilibrium, which increases total welfare from the baseline zero-subsidy policy. In our simple model, if we impose that firms price at marginal cost, or in other words that markets are perfectly competitive, the optimal subsidy goes to zero in both mechanisms. However, competition also comes with a potential drawback: if the subsidy is set too generously, additional competition can actually lead to welfare decreases, as prices may be set below social marginal cost in equilibrium.

To summarize, our simple model highlights the essential points of our empirical investigation that follows. First, increasing subsidies increases consumer surplus but may induce marginal consumers to purchase goods at a social cost higher than their valuations . Second, within the market, subsidies can distort the relative prices of goods, which can lead to allocative distortions. Third, firms can capture some rents from the subsidy by exerting market power and raising prices. Fourth, government expenditures have a social cost due to the deadweight loss of taxation. The social planner's solution balances all of these factors when constructing optimal prices. We note that our model sidesteps an important question of why the government subsidizes markets in the first place. In real world contexts, there may be reasons for subsidization beyond direct within market efficiency considerations, such as consumer-side externalities, redistribution, or political factors. Throughout the paper, we do not take a stand on why there are subsidies in Medicare Part D, but rather take them as given and consider how the present mechanism could be adjusted to obtain more efficient outcomes.

3 Institutional Environment

Medicare is a public health insurance program for the elderly and disabled in the United States that covers over 50 million beneficiaries and costs the government about \$500 billion annually. The program is administered by CMS, and consists of several pieces. Parts A and B cover hospital and outpatient services, respectively, under a fee-for-service model of "traditional" Medicare. Part C, commonly known as Medicare Advantage, was introduced in 1997 and allows consumers to switch from fee-for-service to managed care plans that are administered by private insurers and are highly subsidized by the government. In 2006, Congress expanded Medicare program to include prescription drug coverage via Medicare Part D. In 2016, approximately 41 million individuals benefited from the Medicare Part D program and the Congressional Budget Office estimates that the government currently spends over \$94 billion on Part D annually. This new part of the Medicare program is the institutional setting of our study.

For beneficiaries with traditional Medicare coverage that are not eligible for low-income assistance, buying a Part D prescription drug insurance plans is voluntary and requires an active enrollment decision. These so-called "regular" enrollees may choose one of about 50 stand-alone PDP contracts offered in their geographic region (choice sets in years 2007-2010; in 2016, the average number of choices dropped down to about 20). The Part D program is divided into 34 geographic markets, some of which follow state boundaries, and some bundle states with lower population density. Beneficiaries in traditional Medicare that *are* eligible for low-income assistance are automatically assigned to eligible plans by CMS. These individuals can subsequently change their random assignment by making an active choice. The latter group is known as "LIS choosers." Once enrolled, regular beneficiaries pay premiums on the order of \$500 (see Table 1) a year, as well as deductibles, co-payments or co-insurance. LIS-eligible enrollees receive additional support to cover premiums and cost-sharing.

If either regular or LIS-eligible beneficiaries enroll in private Medicare Advantage plans rather than traditional fee-for-service Medicare, their Part D coverage is provided within the MA plans. This so-called "MA-PD" portion of MA coverage is automatically included in the MA package by the majority of MA plans.

The supply-side of the Part D program has a unique, and controversial, design.⁵ Unlike the rest of Medicare, the drug insurance benefit is administered exclusively by private insurance companies. At the same time, the setting differs from more conventional private insurance markets in two key ways. First, firms are highly regulated and product selection is restricted; CMS sets an annual Standard Defined Benefit (SBD), which defines the minimum actuarial level of insurance that the private plans are required to provide. The SDB has a non-linear structure illustrated in Figure 3; it includes a deductible, a 25 percent co-insurance rate and the infamous "donut hole," which is a gap in coverage at higher spending levels. As long as an actuarial minimum is satisfied, insurers are allowed to adjust and/or top up the SDB contract design, which generates variation in contracts' financial characteristics. In addition, contracts may be differentiated by the quality of insurer's pharmacy networks, which drugs are covered, and other non-pecuniary quality measures.

The second way in which Part D environment differs from more conventional insurance markets is that consumers bear only a fraction of the cost in the program. As much as 90 percent of insurer revenues come from the government's per capita subsidies.⁶ For individuals, who are eligible for low-income-subsidies, these subsidies can go up to 100 percent.

Subsidies are determined through a complex system that depends on firm behavior. First, the government administers an annual "simultaneous bidding" mechanism. According to this mechanism, the insurers that want to participate in the program submit bids for each insurance plan in each region they want to offer. By statute, the bids are supposed to reflect

⁵Oliver, Lee, and Lipton (2004) discusses the political origins of Part D and its mixed reception in the first years of the program, particularly among consumers.

⁶See Table IV.B11 of 2012 Trustees of Medicare Annual Report.

how much revenue the insurer "needs," including a profit margin and fixed cost allowances, to be able to offer the plan to an average risk beneficiary.⁷ CMS takes the bids submitted by insurers for each of their plans and channels them through a function that outputs which part of the bid is paid by consumers in way of premiums, and which part is paid by CMS through subsidies. This function takes the bids of all Part D PDP plans nationwide, adds in the bids submitted by MA-PD plans, weights them by lagged enrollment shares of the plans, and takes the average. Less than 75 percent of this average is set as Medicare's subsidy. The remaining 25 percent of the national bid average together with the difference between the plan's bid and the national average is set as the consumer's premium.

The second feature of the subsidy mechanism concerns the role of low income beneficiaries. CMS utilizes the same insurer bids to determine which insurance plans qualify to enroll randomly assigned LIS beneficiaries. For each geographic market, CMS calculates the average consumer premium. This average constitutes the subsidy amount that low-income beneficiaries receive, known as LIS benchmark or LIPSA. Plans that have premiums below the LIS benchmark qualify for random assignment of LIS enrollees (Decarolis, 2015).

Last but not least, subsidies vary depending on the health of individual enrollees. Insurers receive higher subsidies for sicker beneficiaries through the process known as risk adjustment. Each beneficiary is assigned a continuous risk score that is calculated such that the individual of average health within the Medicare program has a risk score of one. Sicker beneficiaries get assigned higher risk scores. Subsidy payments to insurers are scaled by this risk score to reflect higher expected expenditures that insurers would incur for sicker enrollees. For beneficiaries of especially poor health, insurers also receive a so-called "re-insurance" payment, by which the government pays about 80 percent of individual's spending after this individual has incurred relatively high pharmaceutical costs. This effectively further increases subsidies for sicker beneficiaries.

These mechanisms of risk adjustment and re-insurance are parts of a three-pillar risk equalization system within Medicare Part D. The third part of this system—risk corridors directly decreases insurers' exposure to bottom line risk in profits by capping certain levels of losses (and symmetrically taxing unexpected gains). These three mechanisms play two roles in the market. First, they effectively result in higher subsidies for individuals of worse

⁷There are several nuances buried in the set-up of the bidding procedure that are important for insurers' incentives and will enter the insurers' profit function in our empirical model. First, Medicare sets a minimum required actuarial benefit level that plans have to offer. Plans are allowed to offer more coverage ("enhance" the coverage), but that enhanced portion is not subsidized. Thus, when submitting their bids plans are supposed to only include the costs they expect to incur for the baseline actuarial portion of their benefit. The incremental premium for enhanced coverage has to be directly passed on to consumers.

health. Second, they serve to mute insurers' incentives to cream-skim healthier enrollees by trying to equalize the marginal cost of each enrollee from the insurers' perspective. We discuss how this latter aspect influences our modelling choices on the supply side in Section 4.2.

This somewhat Byzantine system determines prices, and thus allocations, for plans in Part D PDP and, by extension, alternatives such as MA-PD plans that are typically paid the same subsidy. We emphasize three elements of the current system: first, the level of subsidies is endogenous to firm behavior. If all firms increased their bids, the subsidy would be more generous. This generates a complex feedback between firm pricing and the subsidy, as the two are inextricably tied together. In such an environment, it is unclear ex ante whether competition for consumers, which lowers prices, or the softening of demand through the subsidy, which raises prices and the subsidy, will dominate in equilibrium. Second, the mechanism links together the bids of MA-PD plans to the PDP subsidy by including both programs in the calculation of the average bid. This is important because some regions have very low MA-PD bids, so the tie between PDP and MA-PD through the average bid will lead to lower generosity of the PDP subsidy in those markets, which may or may not be efficient depending on why MA-PD bids are low. Third, firms may lower their prices strategically in order to qualify for LIS enrollees; this shading down will also have the effect of lowering the average bid, which decreases the generosity of the subsidy for all beneficiaries. When considered in full, this complex system may or may not lead to socially-desirable outcomes; the remainder of this paper attempts to take first steps in answering this question by understanding the fundamental forces of supply and demand that express themselves through this mechanism in equilibrium. We next turn to describing our model to tackle these questions.

4 Model

We propose an empirical model of demand and supply of insurance contracts in Medicare Part D that will help us evaluate the market structure and the efficiency of the subsidization mechanism in the program. We start with a model of demand for insurance contracts that follows the approach of Berry (1994) and Berry, Levinsohn, and Pakes (1995) (hereafter referred to as BLP). We then move to a supply-side model that allows us to estimate the marginal costs of insurers.

4.1 Demand

We consider two separate demand systems. First, we estimate demand of regular enrollees, who choose their plans, pay full enrollee premiums, and also pay full cost-sharing through deductibles, co-insurance, and co-pays. Second, we estimate a separate demand system for enrollees that are eligible for low income subsidies and face different plan characteristics.

We start with the enrollment decisions of regular beneficiaries. We define the potential market as all Medicare beneficiaries that are not eligible for low income subsidies, and did not receive their Part D coverage through their employer or through special groups like Veteran Affairs. This leaves us with non-LIS Medicare beneficiaries that enroll into a stand-alone prescription drug plan (PDP), or a Medicare Advantage Prescription Drug plan (MA-PD), or do not have any Part D coverage. We let the choice of not enrolling into any part of the Part D program or enrolling through a Medicare Advantage plan comprise the outside option. Within the inside option, individuals are choosing among 40 to 50 stand-alone Prescription Drug Plans (PDPs) that are available in their region. We estimate demand for plans in 34 statutory Part D geographic regions in years 2007 to 2010, for a total of 136 well-defined markets. The utility of the outside option is initially normalized to zero. However, we will consider several counterfactuals where we change the value of the outside option to reflect adjustments to the subsidy levels in the MA-PD program that may happen contemporaneously with changes in the PDP market.

We posit that individuals select insurance contracts among PDP plans by choosing a combination of pecuniary and non-pecuniary plan characteristics that maximizes their indirect utility. We take the characteristics-space approach and project all plans into the same set of characteristics. This approach allows us to make fewer assumptions about how individuals perceive the financial characteristics of plans, but also implies that we remain agnostic about the objective actuarial efficiency of choices, and also do not recover deeper structural parameters such as risk aversion. Despite the fact that we are estimating demand for insurance and thus preferences may depend on risk aversion, we argue that the model of linear index utility with unobserved heterogeneity is suitable for our goals. The risk protection quality of an insurance plan is represented by financial characteristics other than premiums. We can think about the linear utility index as a reduced-form way of capturing revealed valuation of different financial characteristics of plans that are generated by underlying concave utility functions over the distributions of expected spending. In the simulations of the model in Section 6, we will be interested in capturing the demand response to changes in premiums, while keeping the plans' actuarial properties fixed. With these modeling choices in mind, we let the utility consist of a deterministic component and a random shock, ϵ_{ijt} , distributed as a Type I Extreme Value:

$$u_{ijt} = v_{ijt} + \epsilon_{ijt}.\tag{1}$$

The deterministic indirect utility function of a regular enrollee i who chooses plan j in market t is given by:

$$v_{ijt} = -\alpha_i p_{jt} + \beta_i x_{jt} + \xi_{jt}, \tag{2}$$

where p_{jt} is the plan's enrollee premium. Note that unlike in standard product markets, the premium that enrollees pay in Part D is not equivalent to the per capita revenue that firms receive, since there is a large part paid in federal subsidies to insurers. Allowing for the possibility that the government subsidy, z, can be larger than a particular plan's desired per capita revenue, the premium is then equal to $p_{jt} = \max\{0, b_{jt} - z_{jt}\}$, where b_{jt} is the amount of money the firm receives for the enrollee, or its "bid" in Medicare's terminology.

The observable characteristics of plan j in market t, x_{jt} , includes the annual deductible, a flag for whether the plan has coverage in the donut hole, whether the plan has additional coverage beyond the statutory minimum (i.e. it is "enhanced"), and several generosity measures of drug formularies. We also include fixed effects for parent organizations that capture individuals' preferences for brand names of large insurance companies and insurerlevel quality characteristics of plans, such as pharmacy networks, while ξ_{jt} is a plan-specific fixed effect that captures unobserved plan quality.

Several papers, e.g. Ericson (2014), Ho et al. (2015), Polyakova (2016), Wu (2016), Heiss et al. (2016), have documented various behavioral biases related to intertemporal stickiness in plan choice behavior and potential insurers' response to these biases. We take a reduced form approach to this issue by including the vintage of a plan in the utility function as a proxy for inattention and switching costs. The idea is that the longer the plan has been around, the larger the proportion of its enrollees are incumbent consumers from previous years. If there are switching costs or inattention involved in the re-optimization of an insurance plan choice for some consumers, those consumers will appear to be less price sensitive than those choosing a plan for the first time; this reduced elasticity of demand should translate into higher prices on the part of insurers, all else equal. Since we are primarily interested in the effects of the subsidy mechanism on pricing, rather than the effects of switching costs, we do not develop a dynamic model, instead pursuing a reduced form approach of including vintage into the utility function. However, we show in Appendix 8.2 that our approach corresponds

to an explicit structural model of inattention and choice following Hortacsu et al. (2015). A complete characterization of the influence of inattention and switching costs on demand and pricing would require an equilibrium model as in Klemperer (1995) or Dubé et al. (2009). We note that this literature has conflicting predictions about the sign of pricing effects in response to switching costs: Klemperer (1995) concludes that prices are likely to be higher in equilibrium, while Dubé et al. (2009) demonstrate that prices can be lower in equilibrium. Indeed, in the analysis of these issues within the Medicare Part D setting, Ho et al. (2015) and Wu (2016) come to opposite conclusions about which pricing strategies have dominated the market in response to consumer inertia. In either environment, the fact that we are not modelling a dynamic equilibrium may lead to a bias in marginal cost estimates. In the setting of Medicare Advantage plans, Miller (2014) argues that in insurance markets that are characterized by inertial demand, the marginal cost estimates from a static Bertrand model may be around 20 percent higher or lower than the "true" dynamic values. Recognizing this concern in our setting, we report the key counterfactual results in Section 6 for a 20 percent interval around our marginal cost estimates.

Unobserved consumer heterogeneity enters the model through random coefficients on the premium, coverage in the gap, and overall inside option. The unobserved heterogeneity may capture differences in income, as well as individuals' differences in risk and risk aversion. We choose a log-normal distribution for random coefficients on premiums that is only defined on the positive quadrant and reflects typically log-normally distributed income. The random coefficient on premium is composed of a common component, α , and an individual-level random shock, $\nu \sim \mathcal{N}(0, 1)$, which is scaled by σ_{α} :

$$\ln \alpha_i = \alpha + \sigma_\alpha \nu_i. \tag{3}$$

The parameters governing coverage in the gap coverage, β_{gap} and σ_{gap} , and the inside option, β_{inner} and σ_{inner} , are specified analogously but lack the logarithmic transformation.

Estimating demand for LIS choosers is slightly different due to the institutional setting. While LIS enrollees are randomly assigned to eligible plans by the government by default, they are allowed to change their assignment to a plan of their own choice after the random assignment. In general, we model the demand of low-income beneficiaries similarly to what we do for regular enrollees, but there are some differences. The first key difference is that low-income beneficiaries face different characteristics of plans, as their cost-sharing is largely covered by the government. Let the deterministic indirect utility function of a low-income subsidy enrollee i who chooses plan j in market t be given by:

$$v_{ijt} = -\alpha_i^{LIS} p_{jt}^{LIS} + \beta_i^{LIS} x_{jt}^{LIS} + \xi_{jt}^{LIS},$$
(4)

where p_{jt}^{LIS} is the plan's premium for the low-income population. This premium is computed as the remainder of the difference between an insurer's bid and the region-level LIS subsidy (LIPSA), which is higher than the subsidy for regular enrollees. x_{jt}^{LIS} contains observable characteristics of plan j in market t as faced by the low-income population. The difference in plan characteristics for regular and LIS enrollees lies primarily in cost-sharing: to the first order, the LIS population does not face a deductible or coverage in the gap or co-payments above certain thresholds, as this cost-sharing part is picked up by the government.

The second key difference vis-a-vis the regular market is the random assignment of many LIS beneficiaries to plans. The empirical challenge in this setting is that we cannot distinguish between LIS enrollees who are in LIS-eligible plans due to random assignment or by choice. We address this challenge by aggregating all plans eligible for LIS random assignment into one choice within the inside option. To do the aggregation we average the characteristics of these LIS-eligible plans. The idea is to interpret the option of not opting out of the random assignment plans as one distinct choice that LIS enrollees can make. The potential measurement error introduced by this aggregation is alleviated by the fact that plans eligible for LIS random assignment have many of the same key characteristics for the LIS population, such as zero deductibles, no gap in coverage, and otherwise reduced or eliminated cost-sharing. To close LIS demand, we define the outside option for LIS enrollees as choosing an MA-PD plan instead.

4.2 Supply

Modeling the supply side in Medicare Part D market presents a considerable challenge, as the decision-making of the insurers is affected by a complex set of regulatory provisions. We start with a description of the flow of payments in Part D and set-up a general profit function that can incorporate these features. We then discuss our strategy of arriving at an empirically tractable version of the supply-side model.

Firms receive revenues across a variety of channels. For each individual that plan j enrolls, the insurer collects an enrollee premium, p_j . This is augmented with an individual-specific subsidy, z_i , from the government, which is composed of the baseline subsidy for the plan and an adjustment for the enrollees ex-ante health risk, denoted by r_i . For example,

an individual with average risk level will only receive baseline subsidy, while an individual with costly chronic conditions may generate twice the amount of the baseline subsidy in insurers' revenues. This individual-level risk adjustment is intended to make all consumers look equally profitable to firms in order to reduce incentives for risk-based selection. Recalling that the level of the baseline subsidy depends on the average bid, \bar{b} , we can write the subsidy as a function of the average bid and individual-specific health risk as $z_i(\bar{b}, r_i)$. For an averagerisk beneficiary, the sum of the premium and government subsidy is equal to the bid that the firm submitted for that plan.

On the cost side, the *ex post* costs of a plan differ for each enrollee and depend on individual drug expenditures. Some of these costs are mitigated by the government through catastrophic reinsurance provisions, according to which the government directly pays about 80 percent of individual's drug spending for particularly high spenders. Throughout the empirical results we will refer to these reinsurance provisions as "non-premium subsidies" or reinsurance payments. For an individual with a given total annual drug expenditure amount, the costs of the plan will also depend on the cost-sharing characteristics of the plan, denoted by ϕ_j . These include characteristics such as the deductible level, co-pays and co-insurance, as well as coverage in the donut hole if any. We let individual-level *ex post* costs be the function of these cost-sharing characteristics of a plan as well as the individual's measure of health risk, r_i ; that is, we let the cost be $c_{ij}(r_i, \phi_j)$.

The final piece of a plan's *ex post* profit are risk corridor transfers between insurers and the federal government. These transfers that happen at the end of the year, and restrict the downside (but also upside) risk of enrolling extremely costly individuals for the insurers.⁸ We denote the function which adjusts a plan's *ex post* profit with Γ .

The *ex post* profit for plan j as a function of its bid b_j is then:

$$\pi_j(b_j; b_{-j}) = \Gamma\left[\sum_{i \in j} \left(p_j(\overline{b}, b_j) + z_i(\overline{b}, r_i) - c_{ij}(r_i, \phi_j)\right)\right],\tag{5}$$

where the summation is taken over all individuals enrolling in the plan. For each individual i, the subsidy and the cost can be expressed as an individual-specific deviation from the baseline subsidy and an average plan-specific cost of coverage: $z_i = z + \tilde{z}_i$ and $c_{ij} = c_j + \tilde{c}_{ij}$.

⁸See more details in Medicare Part D Manual. As CMS describes in Chapter 9 of Prescription Drug Benefit Manual, risk corridors are: "Specified risk percentages above and below the target amount. For each year, CMS establishes a risk corridor for each Part D plan. Risk corridors will serve to decrease the exposure of plans where allowed costs exceed plan payments for the basic Part D benefit." (See 42 C.F.R, 423.336(a)(2).)

Denote the individual-specific difference in the subsidy and cost as $\eta_{ij} = \tilde{c}_{ij} - \tilde{z}_i$. This function allows us to capture adverse or advantageous selection from the point of view of the insurance plan. Given the empirical evidence in Polyakova (2016) on the selection patterns in Medicare Part D, η_{ij} mostly depends on whether or not a plan offers coverage in the gap. We thus let this individual-specific component be a function of plan characteristics: $\eta_{ij}(\phi_j)$. Using this notation, we can re-write the profit function as:

$$\pi_j(b_j; b_{-j}) = \Gamma\left[\sum_{i \in j} (p_j(\overline{b}, b_j) + z(\overline{b}) - c_j(\overline{r}, \phi_j)) + \left(\sum_{i \in j} \eta_{ij}(\phi_j)\right)\right].$$
(6)

Letting $H_j(\phi) = \sum_{i \in j} \eta_{ij}(\phi_j)$, we obtain a profit function that does not have individualspecific terms and can be written using the market share notation that is useful for the empirical analysis.

As the sum of the premium and the baseline level of the subsidy is by construction equal to the bid submitted by insurer to Medicare, $p_j(\bar{b}, b_j) + z(\bar{b}) = b_j$, we can re-write the pre-risk corridor profit of plan j as:

$$\pi_j(b_j; b_{-j}) = M s_j(b)(b_j - c_j) - H_j(\phi), \tag{7}$$

where $s_j(b)$ is the market share of plan j and M is market size. We emphasize that s_j encapsulates all of the regulatory details involved in turning bids, b, into plan-specific market shares.

We now expand this expression to allow for multi-plan insurance organizations as well as to allow for different prices, marginal costs, and demand from the LIS segment. The structure of profit from LIS enrollment is specified as entirely symmetric to the regular enrollees. We denote quantities related to regular enrollees with superscript R, and quantities related to the LIS part of the market with superscript LIS. The profit function for insurer J offering a portfolio of $j \in J_t$ plans across markets $t \in T$ is:

$$\pi_J(\mathbf{b}) = \sum_{t \in T} \sum_{j \in J_t} \Gamma \left[M_t^R s_{jt}^R(b) (b_{jt} - c_{jt}^R) - H_{jt}^R(\phi) + M_t^{LIS} s_{jt}^{LIS}(b) (b_{jt} - c_{jt}^{LIS}) - H_{jt}^{LIS}(\phi) \right].$$
(8)

Firms maximize profits by choosing bid b for each insurance plan in each market.

Equation 8 is more complex than a standard profit function in a differentiated products market due to how the share equation $s_{jt}(b)$ is constructed. For regular enrollees, the share depends on the plan's premium, p^R , which is not set directly by insurers, but rather depends on the bids of other insurers in a non-linear fashion:

$$p_{jt}^R = \max\left\{0, b_{jt} - \zeta \bar{b}_t\right\},\tag{9}$$

where \bar{b}_t is the enrollment-weighted average bid of all plans in the entire US and ζ is the share of the average bid covered by the federal subsidy. The adjustment ζ is set every year by CMS and is governed by fiscal considerations and the Part D statutes; in 2010, this number was 0.68. The share equation for the low-income segment of the market is substantially more complex. It can be thought about as a piece-wise function with two components: random assignment of low-income enrollees by CMS for those plans that are eligible for random assignment, and enrollment choices by LIS choosers. While LIS choosers are easily modeled in the standard discrete choice demand system, the eligibility requirement for random assignment introduces a discontinuity into the share function. Only plans below the average premium are eligible for random assignment, so for some choices of b_j , the share function for that portion of the market discontinuously jumps to zero.⁹

Section 5.3 outlines how we deal with the discontinuity in the share function in the estimation. Here, we assume that the firm does not anticipate random LIS enrollees and derive the first-order conditions for profit maximization. We also assume that the firm ignores the effect of its bidding behavior on the average bid, \bar{b} ; this seems reasonable in light of the over 1,500 PDP plans that, along with the MA-PD plans, determine the average bid. We assume that firms are small enough so that they take the weighted average as given. Then, for a each plan j in firm i's portfolio of J_i plans, the Nash-Bertrand first-order condition is:

$$\frac{\partial \pi_i}{\partial b_j} = M^R s_j^R(b) + (b_j - c_j^R) M^R \frac{\partial s_j^R(b)}{\partial b_j} + \sum_{k \neq j \in J_i} (b_k - c_k^R) M^R \frac{\partial s_k^R(b)}{\partial b_j} + M^{LIS} s_j^{LIS}(b) + (b_j - c_j^{LIS}) M^{LIS} \frac{\partial s_j^{LIS}(b)}{\partial b_j} + \sum_{k \neq j \in J_i} (b_k - c_k^{LIS}) M^{LIS} \frac{\partial s_k^{LIS}(b)}{\partial b_j}.$$

This expression differs from the more familiar first-order condition in the differentiated product literature in that the market size now plays an important role for the firm's decisionmaking. The market size affects the relative effects on profit from enrolling regular beneficiaries versus LIS choosers. As we now have one equation in two unknowns, marginal costs for regular enrollees, c_j^R , and marginal costs for LIS enrollees, c_j^{LIS} , we make an additional

⁹Decarolis (2015) discusses the piece-wise structure of the share function and the incentives generated by the LIS random assignment mechanism in much greater detail.

assumption to close the model. Following Hsu et al. (2010), we assume that the marginal cost of the LIS enrollees are 13 percent higher than regular enrollees, so that $c_j^R = \kappa c_j^{LIS}$, where $\kappa = 1.13$. Imposing this assumption and collecting terms in vector notation, we arrive at:

$$M^{R}s^{R} - \Omega^{R}(p - mc) + M^{L}s^{L} - \Omega^{L}(p - \kappa mc) = 0.$$
(10)

where

$$\Omega_{rj}^{R} = \begin{cases} -M^{R} \frac{\partial s_{j}^{R}(p)}{\partial p_{r}} & \text{if } \{j, r\} \in F_{f}, \\ 0 & \text{else,} \end{cases}$$
(11)

and

$$\Omega_{rj}^{L} = \begin{cases}
-M^{L} \frac{\partial s_{j}^{L}(p)}{\partial p_{r}} & \text{if } \{j, r\} \in F_{f}, \\
0 & \text{else.}
\end{cases}$$
(12)

4.3 Welfare Metrics

In most of our counterfactual exercises, we will focus on measuring welfare levels and changes for regular enrollees. For these enrollees, total welfare in the Medicare Part D PDP market is comprised of three pieces: consumer surplus (CS), insurer profits (Π), and the deadweight loss associated with taxation used to fund government subsidies (G):

$$W = CS + \Pi - \lambda G,\tag{13}$$

where λ is the social cost of raising revenues to cover government expenditures, G. The welfare metric for the LIS market broadly follows the same idea, although we highlight some differences in the discussion of LIS-specific counterfactuals in Section 6.4.

All three pieces of the welfare function are calculated relative to the outside option. For consumer surplus the normalization to the outside option follows directly from the utility model. In Section 4.1 we had defined utility from enrolling in stand-alone Part D prescription drug plans as being relative to the choice of an MA-PD plan or to not purchasing any Part D coverage. For producer surplus, the insurer pricing decision as formulated in Section 4.2 implicitly takes into account the opportunity cost of serving the outside option. In other words, the marginal cost as recovered from the inversion of the first-order conditions incorporates the opportunity costs of potentially serving each consumer in the MA-PD market or not serving the consumer at all. Consequently, the profit function is defined relative to profits that could have been made in the MA-PD program or elsewhere. Finally, since the government subsidizes both PDP and MA-PD parts of the market, we consider government spending in PDP net of what it would have spent on subsidizing the same individual should they leave the PDP market. We conservatively assume that the outside option for the government is the level of average subsidies spent in MA-PD. This assumption excludes the possibility that some individuals could leave subsidized insurance altogether.

Following Williams (1977) and Small and Rosen (1981), surplus for consumer i with marginal utilities θ_i from plan characteristics, including the premium, takes the following form:

$$CS(\theta_i) = \frac{1}{\alpha_i} \left[\gamma + \ln \left[1 + \sum_{j=1}^J \exp(v_{ij}(\theta_i)) \right] \right], \tag{14}$$

where γ is Euler's constant, and v_{ij} is the deterministic component of utility for person *i* from plan *j* as given in Equation 2.¹⁰ We integrate out over the unobserved taste heterogeneity to obtain consumer surplus:

$$CS = \int CS(\theta) dF(\theta).$$
(15)

The second piece of the welfare calculation is producer surplus that we approximate using the pre-risk-corridor version of the profit in Equation $8.^{11}$

The last piece of net welfare calculations is the deadweight loss associated with raising revenue to cover government transfers to insurance firms and regular Part D beneficiaries enrolled in stand-alone prescription drug plans. In our welfare calculations, we weigh the government spending with the shadow cost of public funds, commonly estimated at $\lambda = 1.3$.¹² Similarly to the outside option reasoning in the case of consumer and producer surplus, we consider how much extra government spending the PDP part of the Part D program (G^{PDP}) generates relative to the outside option of subsidizing the beneficiaries in Medicare Advantage prescription drug plans (G^{MAPD}).

Adding the three parts of the welfare function back together, we have the following

¹⁰Euler's constant is the mean value of the Type I Extreme Value idiosyncratic shock under the standard normalizations in the logit model, and is approximately equal to 0.577. The welfare interpretation of the vintage variable is open to debate; on one hand people may learn to like the plan the longer they are in it, in which case the vintage variable is truly related to utility. On the other hand, if it simply captures inattention, it may not be appropriate to include this as part of welfare. We performed the computation using both interpretations for the key counterfactuals and discuss them in the robustness Section 6.5. Our baseline specification include vintage in the utility function.

¹¹This assumes transfers between firms and the government are welfare neutral.

¹²See, for example, Hausman and Poterba (1987).

measure of average total surplus:

$$W = \int \frac{1}{\alpha} \left(\gamma + \ln \left[1 + \sum_{j=1}^{J} \exp(v_j(\theta)) \right] \right) dF(\theta) + \sum_{j=1}^{J} (b_j - c_j) s_j(p)$$
Social Cost of Government Spending (G)
$$- \lambda \left(\sum_{j=1}^{J} (G_j^{PDP} - G^{MAPD}) s_j(p) \right). \quad (16)$$

While this welfare function describes the surplus for the private market, where firms administer insurance contracts, it does not correspond to the welfare function that a social planner would maximize. If we had the government setting prices for insurance contracts and in effect hypothetically administering these contracts, we would need to take into account the cost of public funds for doing that in the social planner's problem. Another way of thinking about this problem is to imagine that the government dictates prices to private insurers that administer the plans, but then taxpayers cover any shortfall in insurers' profits. Algebraically, both of these interpretations imply that surplus or loss generated in the product market should be weighted with the deadweight loss of taxation. Hence, the social planner's objective function, which we denote with $W^{SP}(p)$ looks the same as equation 16, except that the product market profit is included under the λ -weighted term.¹³ The vector of prices that maximizes this version of the welfare function is the social planner's solution. Note that we use prices in the social planner's case, as the distinction between insurer bids and consumer premiums is not meaningful in this case. Using our notation, optimal prices are (see Appendix 8.1 for algebraic details):

$$p^{SocialPlanner} = \operatorname{argmax} W^{SP}(p) = c + \Delta G + \Omega(p)^{-1} \frac{(1-\lambda)}{\lambda} s(p).$$
(17)

Price is set to marginal cost plus an additional term which adjusts for the opportunity cost of government spending across the inside and the outside options. The final term represents the trade-off between lost consumer surplus and additional product market, a margin which is captured by the social planner.

¹³We also report the outcomes when the profit function is not weighted by λ in Table 11.

5 Data and Results

5.1 Data and Descriptive Facts

Our primary data set combines a variety of aggregate plan-level information released annually by CMS.¹⁴ Table 1 reports several summary statistics of the sample. In years 2007-2010, there were on average 1.3 million Medicare Part D eligible individuals per geographic market in the US. Out of these individuals, on average 0.37 million obtained coverage through their employer or through other sources such as Veteran Affairs. Out of the remaining roughly 1 million on average, about 0.2 million did not purchase any Part D coverage, about 0.25 million chose to buy drug plans bundled with Medicare Advantage, and the rest enrolled in PDP plans. We consider the former two groups as choosing the outside option and focus on demand for PDPs for years 2007-2010.¹⁵

Nationally, consumers in traditional Medicare chose among about 50 PDPs each year in their geographic market. During years 2007-2010, there were on average around 10 large insurer parent organizations in each of the 34 regional markets in which CMS divides the US. Plans are heavily subsidized. While the national average bid was around \$1000, consumers generally paid only a third of that, with the remainder covered by government subsidies.¹⁶ Consumers that chose to opt out of the traditional Medicare program and purchase private Medicare Advantage plans received bundled drug coverage in that program. MA-PD plans received the same types of subsidies as PDP plans for each enrolled consumer.

Panel (a) of Figure 4 illustrates the heterogeneity in pricing across the 34 markets. The graph plots the distribution in premiums by market and plan type, as well as the level of the LIS subsidy. Notably, the plans that offered coverage in the donut hole were consistently more expensive, by a factor of up to three, than plans with only the minimally required coverage. Dispersion in premiums is relatively similar across different markets in the 2010 cross-section. The low-income subsidy averaged \$388, but with substantial variation across markets. Panel (b) illustrates the distribution of market shares across plans within geographic areas in 2010 for regular enrollees. The graph reveals substantial heterogeneity in plan market shares, both within and across markets. While many plans had market share close to zero, some

¹⁴All of the data is publicly available from CMS. CMS tabulates the depository of the data sources at www. cms.gov/Research-Statistics-Data-and-Systems/Research-Statistics-Data-and-Systems.html.

¹⁵Detailed Medicare Part D enrollment numbers are recorded in Tables 14.x of the annual Medicare and Medicaid Statistical Supplement published by CMS at www.cms.gov/ Research-Statistics-Data-and-Systems/Research/MedicareMedicaidStatSupp/Overview.html

¹⁶CMS does not directly report bids; however, they can be backed out from the formula for the subsidization mechanism and the data on premiums and subsidies reported by CMS.

plans covered as many as 20 percent of eligible beneficiaries within a market.

In the remainder of the paper, we repeatedly calculate welfare, both under the observed mechanism along with several simulations of counterfactual market structure and regulatory regimes. The evaluation of welfare requires several assumptions on how we calculate government spending. We calculate government spending on the premium subsidy for standalone prescription drug plans as the sum of per capita (simulated) subsidies multiplied by (counterfactual) plan enrollment. We assume that the average reinsurance, or what we call "non-premium" subsidy for each plan is fixed and does not change across counterfactuals. Thus, total reinsurance subsidies only change across counterfactuals due to enrollment changes. We use CMS annual reporting on average reinsurance payment for each Part D plan as the measurement of the non-premium subsidy.¹⁷ In 2010, for example, the unweighted mean per capita reinsurance payment among PDP plans was \$503 per plan.

In addition to calculating the premium and non-premium subsidies on stand-alone prescription drug plans, we also estimate the government's opportunity cost of having individuals enroll in PDP plans. Based on the choices of individuals switching plans observed in the consumer-level data, we assume that if individuals switch from the inside option of PDP plans to the outside option, they switch to the MA-PD program rather than leave drug insurance altogether. Thus, the government is still likely to incur subsidy spending for these individuals through the MA-PD program. To account for the MA-PD spending, we use CMS data to calculate average observed level of government spending on premium and non-premium subsidies in the MA-PD program. We observe that the average per capita premium subsidy in the MA-PD program is \$686, while the average non-premium subsidy is \$260. This amounts to a total of \$946 government spending per capita on individuals enrolled in the MA-PD program.¹⁸ We use this average spending together with enrollment predictions for inside and outside options to calculate the total opportunity cost for the government of having individuals enroll in PDP rather than MA-PD program. This implies that the per capita government spending is higher on the PDP plans rather than MA-PD plans at the PDP subsidy levels observed in the data, although this relationship can reverse

subsidy levels reported in the text.

¹⁷For raw CMS data, see http://www.cms.gov/Medicare/Medicare-Advantage/Plan-Payment/ Plan-Payment-Data.html.

¹⁸These numbers are calculated using annual Reconciliation Files released by CMS at http: //www.cms.gov/Medicare/Medicare-Advantage/Plan-Payment/Plan-Payment-Data.html?DLSort= 0&DLPage=1&DLSortDir=ascending. For each plan that participated in Medicare Part D, these files report monthly "Average Part D Direct Subsidy" as well as monthly "Average Reinsurance Payment." We take unweighted averages of each of these numbers for MA-PD plans and multiple them by 12 to arrive at the

in counterfactuals.

5.2 Demand Parameters

Table 2 reports the demand estimates. Columns (1)-(3) are for regular enrollees, while columns (4)-(5) are for LIS enrollees. Column (3) reports the estimates of the random coefficient logit model described earlier, which is our preferred specification and the demand system used in the computation of the counterfactual scenarios in the next section. For comparison purposes, we also report estimates for the Berry (1994) logit model; column (1) contains OLS estimates, while column (2) contains 2SLS estimates. Columns (4) and (5) report analogous OLS and IV estimates for the LIS enrollees.

Models (1), (3) and (5) are estimated using instrumental variables. We instrument for plan premiums and assume that other characteristics of the contracts are exogenous in the short run. We motivate this by observing that, while bids for a given plan vary substantially over time, insurers offer a rather stable portfolio of contract types over time (Polyakova, 2016).¹⁹

The concern regarding the bias in the coefficient on premiums is that they might be correlated with an unobserved quality aspect of plans that we fail to capture with the observed characteristics. While we include a rich set of observed plan features, we may not be fully capturing insurer-plan specific customer service or advertising efforts, as well as issues such as drug prices driven by insurers' bargaining power. Some of the variation will be insurerspecific rather than plan specific and so will be captured by insurer fixed-effects. For the remaining variation, we rely on four instrumental variables. Three are BLP-style instruments, measuring the number of PDP or MA-PD contracts that the same insurer offers in the same or different regions. The other instrument is a version of the Hausman instrument measuring prices charged for the similar plans in other geographic markets. Specifically, we construct the instrument by including the lagged enrollment-weighted average of prices of plans offered in other regions in the same macro region and in the other macro-regions by the same company, where macro-regions are defined as three large geographic areas in the US. The idea of this instrument is particularly appealing in our setting due to the regulatory structure of the market, where markets are separated by CMS. Instrumenting the price in one region with the prices of the same contract in other regions, allows us to isolate the variation

¹⁹For example, if an insurer offers a contract with some coverage in the gap in the first years of the program, this insurer is likely to continue offering a contract with some coverage in the gap. The amount of coverage may change, but the dummy-measure that we use for whether there is any coverage in the gap does not appear to respond to short-term demand shocks.

in prices that is common across these contract due to, for example, particular agreements of a given insurer with pharmaceutical producers, and is thus not correlated with marketspecific unobserved quality due to, for example, local marketing. The IV first stage estimates are reported in Table 3. They indicate a positive coefficient for the Hausman instrument, suggesting that plans that cost more in other regions are indeed likely to be priced higher in a given region. The first stage is jointly statistically significant with an F-statistic of 246.

All specifications result in significant and negative coefficients for premiums; as is usual, the estimated price coefficients are biased downward in specifications without instruments. The BLP specification suggests that there is a significant but relatively tight dispersion of the price coefficient in the population.²⁰ To make the comparison between the first three estimates comparable, we note that the estimated modal price coefficient is -12.30, which is close to the IV logit estimate.²¹ The average own-price elasticity for regular enrollees across all plans in Medicare Region 1 in 2010 are -1.19, -4.31, and -5.98, for the OLS logit, IV logit, and BLP specifications, respectively. The elasticity of the BLP specification, which is our preferred estimate, is economically reasonable and similar to the range of elasticities reported in Lucarelli et al. (2012) (-2.0 to -6.0) and Starc and Town (2015) (-5.0 to -6.3).

The remainder of the plan characteristics are estimated to have coefficients with signs according to intuition. Consumers dislike higher plan deductible, but enjoy measures of plan generosity: coverage in the gap, broader coverage of common drugs, and pharmacies in their networks all give higher utility. We also note an economically and statistically significant positive coefficient on the vintage of plans, suggesting that plans that entered earlier in the program were able to capture a larger beneficiary pool. We do not find significant dispersion for two other variables for which we allow random coefficients: the inner option and the dummy for gap coverage.

To give some intuition to the estimated willingness to pay consider the following calculation for the deductible dimension. A typical non-zero deductible varies from \$265 to \$310 dollars. At the average of \$290, removing the deductible has a dollar value of 290 * (-7.11/12.30) = -\$167.63 at the modal value of the premium coefficient. This seems

²⁰We have explored the sensitivity of our demand estimates to a variety of assumptions and specifications. Appendix Section 8.5 reports key coefficients from specifications with alternative sets of random coefficients. We show what happens in the case without random coefficients, as well as when we separately do not have random coefficients on coverage in the gap or the inside option. We find statistical zeros for the random coefficients illustrates that there is a substantial improvement when accounting for unobserved heterogeneity. In all cases we use the sparse grids method as described in Heiss and Winschel (2008) for integration, which dramatically improves the run time of the estimation.

²¹The average own-price coefficient is -13.69 and the median is -13.20.

reasonable, as consumers are not guaranteed to spend through the entire deductible. Specifically, the deductible level lies roughly at the 20th percentile of the spending distribution (Einav et al., 2016). However, removing the deductible does not mean that consumers do not pay anything for the first \$290 in spending. Applying Medicare's standard 25 percent co-insurance rate, we get the expected monetized value of going from \$290 to zero deductible to be 0.8 * (\$290 - 0.25 * \$290) = \$174, which is reasonably close to our willingness to pay estimate. Factor in those consumers that will spend less than the deductible, and \$167 seems like a reasonable estimate.

A similar calculation for coverage in the gap gives a value of $1000 \times 2.89/12.30 = 235$. Again, this is reasonable given that consumers may not enter the gap. Given that even if offered, gap coverage will not cover 100 percent of expenditures, the actual difference between having and not having coverage for the 3000 gap is roughly 75 percent of the gap's value (assuming average co-insurance rate of 25 percent as regulated by Medicare). That is, if individuals spent through the whole gap, they would value the coverage at 2.250. Assuming that consumers that enter the gap have uniform expenses across the gap, the mean gain in coverage is 1125. However, most consumers do not face these costs in the gap; Einav et al. (2016) document about 25 percent of consumers enter the gap. Assuming those consumers entering the gap have uniformly distributed expenditures in the gap, the upper bound on the valuation of coverage is 281. We estimate a value of 235, which is not far from that upper bound.

Columns (4) and (5) report OLS and 2SLS estimates of the Berry logit model for the LIS market. To estimate LIS demand, we shut down the deductible and gap coverage characteristics of plans, as individuals eligible for low-income subsidies receive additional support from the government that helps cover these out of pocket expenditures. We also adjust premiums to reflect the additional subsidies for the LIS enrollees. As described in Section 4.1, we have to make some additional assumptions to formulate a meaningful demand system for the LIS market. The key assumption is that all individuals that we observe in plans that are eligible for LIS random assignment are considered to have chosen one "aggregate" plan, or in other words to have chosen to remain with the random assignment. The results of the OLS and IV specifications for this part of the market are quite similar to the demand estimates for regular enrollees. Individuals prefer plans with more generous formularies and larger pharmacy networks, and plans that have existed on the market for longer time are also more likely to attract beneficiaries. The price coefficient in the OLS specification is similar to the one for regular enrollees. In the IV specification, it is lower, at -7.6, suggesting lower

price sensitivity to prices that, recall, are about \$400 lower per year for the LIS enrollees.

5.3 Marginal Cost Estimates

Following Nevo (2001), we use the demand estimates for regular and LIS enrollees with Equation 10 to recover plan-level marginal costs. As shown in Section 4, the first-order condition linking marginal costs to bids and demand parameters has several non-standard features, particularly the discontinuous share function for LIS enrollees. Therefore, in order to proceed with the estimation of marginal costs, we make several important assumptions. First, we assume that the multitude of risk adjustment and re-insurance mechanisms implemented in Medicare Part D imply that insurers de facto face constant expected marginal costs. Second, we select a subset of plans that were plausibly not distorted by LIS gaming.

In essence, the idea is to select a group of plans for which we find the Bertrand-Nash assumption acceptable for describing the pricing behavior of the insurer. We construct a group of such plans by selecting all contracts of those insurers that within a given market (year-region) were not eligible to enroll randomly assigned LIS individuals into any of their plans. Even if we assume that this group of "non-manipulating" plans does not have premium distortions related to the LIS benchmark, we may be still be worried that these plans are not comparable to plans qualifying for low-income enrollees. Empirically, this does not seem to be the case. There has been substantial variation in the low-income subsidy across regions and there are many insurers who never qualified for low-income enrollees in at least one region. This variation is mostly due to the different penetration of Medicare Advantage: where in 2006 enrollment in Medicare Advantage was high, MA-PD received a high weight in the calculation of the low-income subsidy and, since their premium is typically close to zero, they induced a small low-income subsidy (Decarolis, 2015).²² With this in mind, we proceed to invert the first-order conditions only for these "non-distorted" plans. Since the plans offered by the same insurer across different regions are remarkably similar, the marginal cost estimates of the "non-distorted" plans through the inversion of the first-order condition can be used to predict the cost of similar plans in other regions for which we could not directly infer costs.

In the first step, we recover the marginal costs for "non-distorted" plans using Equation 10. We use 756 plans offered in 2010 data, as our counterfactuals will focus on this year only. We then proceed to relate the estimated marginal costs to the observed characteristics

 $^{^{22}}$ The variation in the total weight assigned to MA-PD in 2006 is substantial ranging from almost 60 percent in Arizona and Nevada to less than 4 percent in Mississippi and Maine.

of non-manipulating plans by estimating the following hedonic-style linear regression:

$$mc_{jt} = X_{jt}\beta + \delta_t + \tau_j + \epsilon_{jt},\tag{18}$$

where X_{jt} includes the same non-premium characteristics of plans that we had included in the utility function. We add the unobserved quality estimate for each plan as an additional explanatory variable in X. We condition the regression on firm and market fixed effects to account for inherent differences in marginal costs across insurers and geographic regions.

Table 4 reports the coefficients for the hedonic regression. We note that the most important determinants of marginal costs appear to be, as expected, the plans' coverage limits, as well as the generosity of their drug formularies. For example, we estimate that offering coverage in the gap increases a plan's marginal cost by \$422 a year, which is a large increase relative to the average marginal cost estimate of about \$1,171 from the inversion procedure. This estimate of the additional marginal cost from coverage in the gap roughly corresponds to the premium add-ons that are charged by insurers that offer coverage in the gap.

We use the estimates of how plan characteristics translate into marginal costs to predict marginal costs for all plans that we did not include in the inversion procedure. This exercise hinges on the assumption that all plans have a similar "production function." In other words, we assume that the plans that manipulate the LIS threshold manipulate the premiums, but do not have different marginal costs conditional on a set of non-price characteristics. This appears reasonable, as the main source of costs in the insurance market is determined by individual health risk; therefore, it is conceivable to assume that plans with the same financial characteristics and formulary generosity will have similar marginal costs.

Figure 5 plots the distribution of predicted marginal costs and compares it to the estimated distribution via the inversion procedure. We estimate substantial heterogeneity in the marginal costs across plans. We find this heterogeneity both through the inversion procedure and in the hedonic projection exercise. This indicates that our hedonic-style regression captures the key drivers for the differences in marginal costs. The manipulating plans are estimated to have slightly lower marginal costs on average, which is intuitive if we believe that cheaper plans are the ones that would try to compete for LIS enrollment. The marginal cost estimates in both cases are centered around \$1,100 and range from about \$750 to about \$1,900.

Our marginal cost estimates imply mark-ups in the order of 9 percent on average. The bottom panel of Figure 5 plots the full distribution across plans. These mark-ups are fairly low, suggesting that the environment is reasonably competitive. We explored several external

sources that discuss the profitability of the Part D PDP plans to verify that our estimates of marginal costs and mark-ups appear plausible. A CBO report from July 2014 (Congressional Budget Office, 2014) provides, to our knowledge, the most detailed publicly available external analysis of revenues and costs in stand-alone prescription drug plans for years 2007-2010. According to Table 2-1 of this report, in 2010, PDP stand-alone plans collected on average \$1,136 in direct subsidies and premiums per person. This measure is similar to our measure of bid in the mark-up formula (bid-marginal cost)/(marginal cost). To construct a proxy for marginal cost from the CBO report, we take the estimate of the "Net Drug Spending for the Basic Benefit," which is reported to be \$1,382 and subtract the reinsurance payments that are reported to be \$521 on average per person.²³ This gives us an estimate of direct costs for basic benefit of \$861. Using the latter number—which certainly is within the scope of our estimated marginal costs—and the average revenue reported above, we would arrive at a mark-up of (\$1,136-\$861)/\$861, or 32 percent. These numbers, however, do not include administrative costs (which CBO reports to be \$242, but it is calculated as a residual and includes profits and administrative costs), and further does not include "enhanced" benefits that are covered by additional premiums and are included in our calculations. To arrive at around a 10 percent markup, we would have to assume that administrative costs are around 170 on average per person: (1,136-861-170)/(8861+170)=10 percent, or 70 percent of the number that CBO reports as an estimate of administrative costs plus profits, which appears plausible. Overall, our estimates appear realistic and consistent with the actual cost data.

5.4 Welfare Estimates of the Current Mechanism

Using the demand and marginal cost estimates from above, we next compute consumer surplus, producer profits, government transfers, and total surplus for the observed market allocation, as outlined Equation 16. For expositional clarity, we report results for a single year (2010). The calculations are reported in the first and second columns of Table 5 for regular and LIS enrollees, respectively.

We estimate that total annual consumer surplus generated by Part D PDP for regular enrollees was about \$2.5 billion, or about \$280 per enrollee.²⁴ We also estimate that insurer

 $^{^{23}}$ Since our marginal cost estimates come from the inversion of the insurers' first-order conditions, they are net of reinsurance transfers between the plans and the government.

²⁴This is calculated as the sum across markets of average consumer surplus in each market multiplied by the market size of each market. Note that the consumer surplus estimate relies solely on the demand model for regular enrollees and does not depend on any assumptions or specification of the supply side.

profits, before risk corridor adjustments, were about \$450 million.²⁵ In sum, the Part D PDP market generated a little less than \$3 billion of consumer and producer surplus. This surplus came at a price, however, as CMS reports that the government spent nearly \$9.5 billion in subsidies in Part D PDP, split between \$6 billion on premium subsidies and \$3.5 billion on non-premium subsidies. Under our assumption that the government would still have to pay subsidies if individuals were to leave PDP and switch to MA-PD, we also estimate the government would have spent \$8.2 billion on the same individuals were enrolled in MA-PD. The difference between the total PDP subsidies and what these would have been in MA-PD, a little more than \$1 billion, is the opportunity cost of government spending on PDPs, which, along with our assumption that the deadweight loss of government taxation is 30 cents on the dollar, gives us the social cost of government spending of \$1.4 billion.

We estimate the total surplus generated by the regular Part D PDP market was about \$1.5 billion. This is one of the primary findings of the paper, along with the observation that the vast majority of this surplus came from foregone government expenditures. That finding follows directly from our assumption that all enrollees in Part D PDP would have been enrolled in MA-PD had the Part D PDP market not existed. This is likely an overestimate of government spending, as it is possible that some of the enrollees in Part D PDP would have decided not to obtain drug coverage of any variety. However, we think that this assumption provides a plausible upper bound for the welfare gains of the Part D PDP program. We also compute a lower bound by assuming that the opportunity cost of government spending was zero. In that case, we estimate that a dollar of public funds would have generated only 0.17 dollars of surplus per subsidy dollar. That is, welfare would have been improved dramatically by shutting down the program entirely. However, if we do take the government's payments for the outside option into account, we arrive at a more encouraging calculation of 1.11 dollars per opportunity-cost-dollar spent by the government. The break-even opportunity cost of government spending is \$261 per enrollee; if government spending outside Part D

²⁵In the estimation of insurer profits from regular enrollees in the counterfactuals, we let the per capita revenue from regular enrollees be equal to the bid of the plan as simulated in the counterfactual plus the enhanced component of the premium as observed in the data if the plan is enhanced. We thus assume that the enhanced component of the premium does not change across counterfactual mechanisms and does not enter bidding. We then take the difference in estimated marginal cost and counterfactual per capita revenue for each plan and multiply it by the counterfactual share of each plan in each market scaled by the market size of the regular enrollees' market. Relative to the profit function formulation in Equation 8 we are not explicitly calculating the effects of risk-corridors that may alter profits at the end of the fiscal year. Our welfare estimates would be negatively affected if risk corridor payments to firms were positive, as the transfer would be adjusted by the deadweight loss of raising taxes. On the other hand, transfers from firms to the government would be welfare neutral. We are further not explicitly calculating the selection component of the profit function $H(\phi)$, which is zero in expectation by construction.

PDP exceeds this number, the Part D PDP market generates non-negative social surplus.

Consumers in the LIS market enjoyed \$2.8 billion in surplus, driven higher in part by lower cost-sharing in the LIS program. We do not report the profits associated with this part of the market. The reason is that application of the static Bertrand-Nash model of competition used to recover marginal costs does not apply to the dynamic competition firms are engaged in for LIS enrollees. Computing government subsidies and government opportunity cost for LIS enrollees requires some additional accounting, so as to incorporate LIS-specific payments to insurers that cover the generous reductions in cost-sharing that LIS beneficiaries enjoy. We add LIS premium subsidies to the row that counts government premium subsidies in PDP. For the non-premium subsidies, we add the per-plan average payments for LIS cost-sharing that CMS reports for 2010. These generate significant quantitative changes to subsidy levels as compared to regular beneficiaries — the average of the LIS-specific non-premium subsidy in 2010 alone was \$626. We do similar accounting adjustments on the MA-PD side, so as to make the opportunity cost calculation comparable to the calculation of PDP subsidies. After these adjustments and under insurer profits set to zero, total welfare is computed at \$3.1 billion. The fact that this number is twice as high as in the regular market, despite similar inside option enrollment level, is driven by two factors. First, LIS beneficiaries face much lower prices — on average \$480 lower annual premiums than regular enrollees. Second, given the numbers on the level of subsidization of LIS enrollees in both PDP and MA-PD programs that we were able to obtain, both programs appear roughly equally expensive. Hence, the net government spending component of the welfare function for LIS enrollees is close to zero, which implies that total surplus is closer to the consumer surplus number than it is the case in the regular market, where PDP enrollment is expensive from the government's perspective.

While the total amount of surplus generated in both markets is positive, this calculation does not provide any guidance about how efficient the current subsidy mechanism is relative to the first-best, or how its various components contribute to the overall welfare outcomes. To address those questions, we next turn to assessing the welfare properties of counterfactual mechanisms which change the way that subsidies are determined.

6 Counterfactual Results

We are interested in understanding how insurer incentives, the subsidy mechanism, and their interplay determine market outcomes. Recalling the framework in Section 2, the essential
economic ingredients are consumer willingness-to-pay, the exercise of market power, and the generosity and form of the subsidy mechanism.

6.1 Links Across Markets

To isolate the effect of the current subsidy mechanism for regular enrollees, we begin by simulating the market once we remove the interconnections existing in the subsidy formula with features of the LIS and MA-PD portions of the market. In the first counterfactual, we remove LIS enrollees from insurers' profit functions while holding all other aspects of the current subsidy mechanism unchanged. Hence, prices are set as we describe above in Section 4, where the baseline subsidy is a fraction of the enrollment-weighted average bids from Part D PDP and MA-PD. In the second counterfactual, we remove MA-PD bids from the computation of the weighted average bid. In this case, the baseline subsidy is set as 68 percent of the average Part D PDP bid. This second counterfactual gives us a benchmark to which we can compare other subsidy mechanisms for regular beneficiaries, keeping separate the issues associated with LIS assignments and MA-PD bids. The outcomes of these two counterfactuals are described in Columns (3) and (4) of Table 5.

The results of the first counterfactual where the tie between regular and LIS enrollees is severed illustrates the rich equilibrium effects of changes to the subsidy mechanism. The key change is that firms previously had incentives to lower prices of some of their plans in order to qualify for random LIS enrollments. This incentive would lower average prices, which in turn lowers the generosity of the baseline subsidy. Removing this incentive leads to two related equilibrium effects: firms have incentives to raise bids since they are no longer competing to get under the average premium, and the baseline subsidy increases as the average bid increases. There is a feedback effect between the two forces, as higher bids lead to more generous subsidies which in turn lead to higher bids and so on; we compute the equilibrium that results when firms' projection of the average subsidy is equal to the one resulting in equilibrium under those beliefs. This is born out in the results in column (3); the average government premium subsidy per enrollee increases from \$680 to \$739. The combination of the more generous subsidy, along with a shift in the relative prices of Part D PDP plans, results in slightly lower enrollment-weighted average premiums (\$473 vs \$502 in the data), and a seven percentage point increase in the share of the inside option. Overall consumer surplus increases, but tellingly the average consumer surplus declines from \$287 to \$274 per enrollee. The role of the LIS incentives in disciplining firm prices is shown in the dramatic increase in firm profits, which more than double to over \$1 billion. Echoing the discussion in Section 2, increasing the generosity of subsidy induces relatively low marginal valuation consumers to enter the market, with some of the rents being captured by firms. The question becomes whether the extra value generated by the change in the mechanism exceeds the social cost of providing the goods. In this case, the answer is a clear no: although consumer surplus and producer profit increase by almost \$1 billion, government expenditures increase by more than \$2.4 billion. Accounting for the opportunity cost of government spending, total welfare declines by \$44 million when the regular and LIS enrollee markets are unbundled. This counterfactual mechanism results in an inefficient system: the per dollar return on government spending drops to about 13 cents per nominal dollar and only 62 cents per opportunity cost dollar.

The removal of the MA-PD plans from the enrollment-weighted average bid used to compute the baseline subsidy results in even stronger effects along the same lines. As the MA-PD bids are quite low and thus typically drive down the average bid, removing their bids increases the average bid both mechanically and due to insurers' strategic response. Consequently, the premium subsidy further increases to \$782 per enrollee, with the average premium declining to \$420. Enrollment in the market increases to 55 percent and producer profits increase by more than \$200 million. While total consumer surplus increases by another \$520 million, the marginal consumers drop the average consumer surplus per enrollee to \$263. Higher subsidies also imply higher government spending of \$15 billion on Part D PDP subsidies, and thus the net result is a decrease in total surplus of \$540 million. The combination of these factors drives the per dollar return on government spending to about 6 cents per nominal dollar and 26 cents per opportunity cost dollar.

The results of these counterfactuals are instructive. They highlight the complex equilibrium interactions arising here between the endogenously-determined subsidies and social welfare along the channels discussed in Section 2. However, they also illustrate the potential for large welfare effects arising from seemingly-unrelated choices in mechanism design. While it is possible that the designers of the Part D subsidy mechanism were able to predict the directional effects of tying prices in the Part D PDP market for regular enrollees to the eligibility of those same plans in the LIS market, and of linking the baseline PDP subsidy to MA-PD bids, it seems unlikely that anyone could *ex ante* predict the exact magnitudes of those choices. As it turns out here, the designers' choice to link on both of those margins resulted in higher levels of welfare. However, we emphasize that, again, the welfare increase only comes from the foregone opportunity cost of government expenditures in Part D PDP. Both consumer surplus and producer profit are higher when the links are removed. The policy implication revolves solely around the ability of the mechanism to prevent socially-costly government expenditures. That does not have to be the case in general; as we show below, there are a range of alternative counterfactual mechanisms that may do better on all fronts. First, we explore the role that market power plays before considering alternative subsidy mechanisms.

6.2 Market Power

Before proceeding to counterfactuals that change the subsidy mechanisms, we first evaluate the effect of market power. We consider the current CMS mechanism (without the LIS and MA-PD pricing links) and explore two polar cases: full competition and monopoly. We implement these two counterfactuals by simulating alternative plan ownership structures within each CMS-defined market. In the first, we assume that each PDP is its own firm; in the second, we assume that every PDP in each market belongs to one firm. The results are reported in the last two columns of Table 5. Our results are somewhat surprising on first glance: contrary to the usual introductory microeconomics intuition, social welfare basically remains unchanged under atomistic competition but increases substantially under monopoly ownership. Under atomistic competition, consumer surplus and producer profits actually decline modestly, driven by the fall in premium subsidies per person to \$771. The net effect of the three changes is essentially zero. However, under monopoly consumer surplus decreases by \$600 million to \$2.8 billion and producer profits go up almost \$900 million to \$2.2 billion. The real key to increasing social welfare is that the market share of the inside option declines to 45 percent; the consumers who remain in the market are those with relatively high valuations. The combination of these factors pushes welfare higher to \$1.2 billion under monopoly. This comports again with the idea that, holding the subsidy mechanism fixed, market power may actually be beneficial in preventing consumers with valuations lower than social cost from entering the market.

The results suggest that, without links to the LIS market and MA-PD, the current ownership configuration delivers outcomes fairly close to that of a purely competitive ownership structure. This is an interesting result, as one of the motivating reasons for using managed competition to deliver publicly-subsidized goods and services was to leverage competition to reduce prices. Our conclusion on this point requires caution, however, as we cannot assess the counterfactual of possible alternatives, which may include a standard government-run program or a regulated monopoly. For one, we do not know anything about the government's, or regulated firms', comparative marginal costs of delivering similar insurance plans in such a scenario, and therefore we cannot make any claims about additional efficiencies introduced by competition. Second, we take the marginal costs of firms as given here; it is likely that a single buyer would be able to exert monopsony power in negotiating with upstream pharmaceutical companies. The combination of the two effects is ambiguous, and we remain conservative in our conclusion that the current ownership structure gives results similar to that if all plans were independent firms.

6.3 Alternative Subsidy Mechanisms

So far, we have considered the effects of linking the current subsidy mechanism across markets, and the role of market power within that mechanism. We now start to explore alternative subsidy mechanisms in Part D PDP. Our goal is two-fold: one, to understand how supply, demand, and the subsidy mechanism interact by manipulating various components of the economic environment; and two, to show how total surplus changes across various mechanism choices.

One issue is that our baseline modeling framework holds the world outside the Part D PDP market constant. This is well-suited for assessing the additional surplus provided by the market, but once we consider counterfactual environments it is natural to ask why the government would restrict changes to only the Part D PDP market. To address the concern that the government would change the mechanism in both Part D PDP and MA-PD markets simultaneously, we consider two classes of counterfactuals which differ by how we treat the outside option. In our baseline set of counterfactuals, we hold the value of the outside option fixed. In a second set of counterfactuals, we adjust the value of the outside option to reflect the possibility that changes in the Part D PDP market would also induce changes in the MA-PD market. We implement the changes in MA-PD by calculating the counterfactual subsidy in the Part D PDP market in the baseline counterfactual and applying this counterfactual subsidy to observed MA-PD bids. We then calculate what this change in MA-PD subsidy implies for changes in average MA-PD premiums, which gives us the change in value of the outside option. This adjustment is not perfect, as it assumes that the average value of the plans in the outside option move in lockstep with the inside option, but we intend for it to capture the broad implications of making the inside and option both much more (or less) generous.²⁶

²⁶We discuss the details of the adjustment calculation in Appendix 8.4.

6.3.1 Counterfactuals: Outside Option Fixed

We start with the counterfactuals without adjustment to the outside option. Some of these counterfactuals are quite extreme, such as requiring consumers to pay radically higher premiums than observed in reality, but we emphasize that the point of those more extreme exercises are not to consider the computed numbers literally; rather, we want to emphasize that these extreme counterfactuals are a method for understanding the economics of the marketplace. By pushing one or more of the policy levers to an extreme, one can highlight the changes that occur across various levels in the system, helping deepen our understanding of the market's workings.

Keeping the outside option fixed, we start by asking whether deviations in the subsidy mechanism form the averaging rule currently used to proportional or lump-sum subsidies could improve total welfare.

Proportional Subsidies We start with a proportional subsidy mechanism in which premiums are given by: $p_{jt} = xb_{jt}$, with $x \in [0, 1]$. Table 6 the findings for three levels of x: 5 percent, 32 percent, and 95 percent. The idea of the 5 percent counterfactual is to test how insurers would respond if the government almost entirely bore the cost of purchase. With the almost-complete reduction in consumer price sensitivity, insurers dramatically increase their bids: we find a greater than threefold increase in bids. Consequently, even though individuals now pay only 5 percent of the bids, the premiums are still relatively high at \$190 a year on average. This drop in premiums, however, is sufficient to increase enrollment to 100 percent in the PDP plans from 55 percent in the baseline. Consumer surplus increases to \$13 billion, and insurers are able to capture huge amounts of the rents generated by the subsidy—profits shoot to over \$50 billion. Given the dramatic increase in bids and enrollment, government spending increases enormously. The result is a stark drop in welfare levels to an astounding negative \$37 billion. Similar intuition prevails in the less extreme counterfactual at 32 percent, which corresponds to the premium fraction that CMS currently uses after taking the weighted average of bids to determine premiums and subsidies. In both cases, welfare losses are driven by low-valuation consumers entering the market, the subsidy being far too generous, and firms capturing most of the rents. The two more generous proportional subsidies act as extremely expensive mechanisms for transferring funds from taxpayers to firms.

A different story emerges when consumers have to pay 95 percent of the bid; effective consumer prices nearly triple over the baseline level. At this level of prices, very few individuals, about 1 percent, are willing to purchase PDP plans and switch to the outside option. Despite the drop in enrollment, the program generates even more surplus than the benchmark case, and has a very high surplus per subsidy dollar. Such a high surplus per dollar is not surprising, as the government pays little and only the beneficiaries with the highest willingness to pay participate in the program.

This case highlights one of our central empirical findings: all else equal, there is little innate willingness to pay for the Part D PDP plans outside of their subsidies. This is likely due to the fact that there are close substitutes readily available at very low (even zero) prices in the form of MA-PD. While traditional Medicare and Medicare Advantage are not identical substitutes, the 95 percent counterfactual highlights that consumers are willing to switch to MA coverage if they face full price of pharmaceutical coverage with traditional Medicare. Of course, we have kept the subsidies in the MA-PD program unmoved, which makes the outside option look much more attractive. We return to this point below when we re-introduce a link between the generosity of the subsidy in PDP and MA-PD.

Overall, these results can be interpreted by considering that proportional subsidies have two effects relative to the observed mechanism. First, there is a price level effect, as a very generous subsidy would decrease the overall level of premiums faced by consumers, even though firms capture much of these rents by increasing their bids. Second, there is a significant change in relative prices that makes the more generous plans relatively more attractive. The counterfactuals illustrate the strong impact of subsidy structure on insurer behavior. In cases where consumers do not face 100 percent of the extra premium in more expensive plans, competitive forces are significantly muted and insurers pass through substantially higher expenditures to the inelastic federal budget.

Flat Subsidies (Vouchers) The second alternative subsidy system that we explore is a flat voucher mechanism, where F is given to each Part D-eligible consumer to be spent on a Part D PDP plan. The premiums are set to equal to: $p_{jt} = \max\{0, b_{jt} - F\}$. The last three columns of Table 6 report results for three flat subsidy levels: \$0, \$700 and \$1400.

Column (4) shows that F = \$0 is associated with such high premiums that essentially nobody enrolls in PDPs. Not surprisingly, this case is very similar to that in the preceding column of Table 6. At the other extreme, F = \$1400, which is double the subsidy level observed in reality, turns out to be a sufficiently generous subsidy to guarantee 100 percent enrollment in PDPs. While full enrollment in Part D PDP plans was also attained with the 95 percent proportional subsidy, the total welfare in the two cases is remarkably different. Although both entail negative welfare, the flat subsidy does far better: a loss of \\$9.8 billion relative to \$37.4 billion. The key difference is that consumers still must pay every marginal dollar for plans priced above the voucher, so while insurers utilize the increase in subsidy and increase their bids, they do not do so dramatically. In fact, almost all firms submit bids at exactly the voucher. Consumers face essentially zero premiums, leading to an increase in consumer surplus to \$13 billion, which is essentially the upper bound without resorting to giving consumers money directly. Producer profit increases substantially to \$4.8 billion. However, as seen previously, the social cost of government expenditures overwhelm these welfare gains, resulting in a large net welfare loss.

The third case, F = \$700 is particularly noteworthy. Although the generosity of the voucher is nominally close to the observed subsidy, \$676, incentives are very different under a flat voucher relative to the observed mechanism. While in the observed mechanism, the subsidy is determined after the bidding process as a fraction of the average bid, here the subsidy is set *ex ante* and it does not respond to endogenous insurer behavior. Most importantly, the voucher requires enrollees to pay dollar-for-dollar for premiums above the voucher level, preserving a substantial demand elasticity and helping discipline equilibrium bids. Compared to the benchmark case, we find that prices that consumers face increase by about \$100. This leads to a drop in enrollment to 37 percent, compared to the benchmark 55 percent, and an accompanying decrease in total consumer surplus and producer profit. Government spending patterns, however, change starkly in response to decreased enrollment and thus overall, the total surplus under the voucher, \$1.8 billion, is higher than in the benchmark case of \$970 million.

The change in welfare across the range of vouchers suggests that there is an optimal voucher above zero. Therefore, we repeat the calculation of total welfare at a range of vouchers from \$0 to \$1500 in increments of \$100. Figure 6 summarizes the outcome of these calculations graphically. We find that the total welfare is the highest at F =\$600. Setting vouchers higher than that level significantly reduces welfare. Less generous vouchers lead to lower welfare, but the gradient is not as steep and thus the cost of deviation from the optimum is lower. Setting the PDP voucher at zero, still leads to total welfare of \$1.1 billion. This is \$900 million less in total surplus than at the optimal voucher. Setting the voucher at \$600 above the optimum, on the other hand, results in a welfare loss of about \$4 billion. Column (5) in Table 7 reports the welfare estimates for the optimal national voucher. Total surplus is \$2 billion, or twice the baseline case.

In looking at welfare outcomes at different voucher levels, we held the voucher constant at the national level. At the same time, we have documented substantial heterogeneity in demand, supply, and prices across the 34 geographic markets. Thus, the next dimension of regulatory intervention we explore is allowing the government to set geographically-differentiated subsidies across regions. We focus on the flat voucher mechanism, as regionally-differentiated vouchers would be the simplest policy change to implement. To implement this counterfactual, we compute welfare at different levels of possible vouchers (from \$0 to \$1500 at \$100 steps) within each region, and then for each market select the voucher that results in the highest welfare within that region. We find that in 19 out of 34 markets, it is optimal to set the same voucher subsidy that would have been the uniform optimum of \$600. In other regions, however, it would be welfare-maximizing to deviate from this subsidy. We find that in two markets, it would be optimal to offer higher subsidies of \$700 and \$800, while in the remaining markets it would be welfare-increasing to lower subsidies by \$100-\$200. Column (6) of Table 7 reports the welfare calculations for these market-specific optimal vouchers; surprisingly allowing for geographically-differentiated vouchers would only result in a small increase in total welfare of \$40 million to \$2.1 billion.

Direct Mechanisms The relatively good performance of the voucher system compared to both the current system and the proportional subsidy indicates a potentially simple policy reform that could improve welfare. Nevertheless, it is interesting to quantify to what extent such a decentralized approach manages to move the market closer to some ideal benchmark. In this final part of the analysis, we thus compare the level of welfare attainable under marginal cost pricing and by the social planner. Columns (1)–(3) in Table 7 reports the full set of welfare estimates for private marginal cost, social marginal cost, and the social planner, respectively.

We start by setting premiums equal to the estimated marginal costs for each insurance plan: $p_{jt} = MC_{jt}$. The results are reported in column (1). Facing premiums as high as marginal costs, consumers leave the PDP program in favor of the outside option with enrollment dropping to 1 percent. The level of total surplus is nevertheless high, since only consumers with the highest willingness to pay enroll in the program, while the government is paying relatively little as it provides only "non-premium" subsidies.²⁷ We expand upon this counterfactual and let consumers face the full social marginal cost rather than only the marginal cost of insurers that was estimated under the existing reinsurance subsidies. To

²⁷The details of how we calculate "non-premium" or reinsurance subsidies are discussed in Section 5. Recall that these represent government payments in the catastrophic coverage part of Part D plans, whereby for the individuals with the highest spending levels the government directly pays 80 percent of expenditures after a certain threshold. This effectively amounts to additional subsidy dollars that, however, do not directly depend on insurers' bids.

calculate the premiums, we add average observed re-insurance subsidies (RIS_{jt}) for each plan to the estimated marginal costs: $p_{jt} = MC_{jt} + RIS_{jt}$. The outcomes, reported in Column (2) of Table 7, are similar, albeit starker. In this case, enrollment and total surplus fall even more, with nearly none of the market choosing the inside option of PDP plans.

To put an upper bound on possible surplus, we solve the social planner's problem. As expected, the social planner's problem generates the highest total surplus at \$3.6 billion. This surplus comes at a cost of large subsidies. We calculate that the optimal prices in PDPs are on average lower than in the benchmark case at \$377. In addition to premium subsidies, the government carries the full cost of the program, including the coverage of insurer losses of nearly \$10 billion. As the algebraic expression for social planer's prices in Equation 26 suggests, the social planner sets prices for each plans as a function of this plan's social marginal cost and a fiscal adjustment term. The latter takes into account how much enrollment in a given plan would cost the government.

Contrasted with the optimal voucher results reported in the last two columns of Table 7, voucher welfare is higher than that generated under (social) marginal cost pricing. However, not even the market-specific optimal voucher manages to get near the welfare value under the social planner. The social planner's solution demonstrates the idea that at flat subsidy rates that are unrelated to the efficiency of individual plans, subsidies distort the allocation of individuals across plans within a given market. Figure 8 shows this point graphically on the example of California's market in 2010. Relative to observed prices, the social planner's solution is to increase premiums in plans with higher social marginal costs. This results in the re-allocation of individuals across plans—market shares of plans with lower social marginal cost increase, while the market shares of plans with higher social marginal costs decrease. Note, however, that the effects are not monotonic. For example, for some plans with coverage in the gap, which have high social marginal cost, prices increase substantially, but the market share almost doesn't change, as there is still enough willingness to pay for at least some plans with generous coverage in the market. The message of the social planner's solution is clear: consumers in the market are systematically choosing plans that are too socially expensive. The social planner ameliorates this shortcoming by increasing subsidies on less-costly plans and increasing premiums for the more expensive plans.

Finally, we consider an alternative scenario that captures some of the intuition of the social planner while retaining simplicity. In a counterfactual that simulates the idea of a very generous public option (which can be thought of as, for example, expanding traditional Medicare to have pharmaceutical coverage), we replace all plans in each market with the

characteristics of the plan with the lowest estimated marginal cost.²⁸ The government then pays the same subsidy as under the observed allocation, i.e. \$676. The results are reported in Column (4) of Table 7. In this counterfactual simulating public option coverage, we find that overall surplus is lower than that under the social planner, but is very high relative to other counterfactuals, at \$2.8 billion. Consumer surplus is very high while producer profit is set to zero by fiat. This counterfactual manages to achieve a near-optimal sorting of consumers between MA-PD and PDP, leading to the best outcome across all counterfactuals save for the social planner. This counterfactual is also particularly appealing given its simplicity in theory; in practice, it of course depends on the ability of the government to offer a public plan at the cost of the cheapest private plan observed in this competitive environment. One way in which this kind of semi-public option could arise would be through an auction mechanism, where only one — most efficient — private plan is allowed to serve a geographic market for a given year.

6.3.2 Counterfactuals: Adjusted Outside Option

We now turn to looking at the counterfactuals where the outside option is adjusted to reflect changes in the generosity of subsidies in Part D PDP. Using the changes in the baseline counterfactuals, we repeat those counterfactual calculations with the outside option adjusted by the average value of the change in Part D PDP generosity.²⁹ We intend for this adjustment to capture the idea that CMS may adjust subsidy generosity in the MA-PD market to mirror changes in Part D PDP. The results are given in Tables 8 and 9. There are two additional rows included in the tables, listing the additional payments that the government makes to the MA-PD in each counterfactual relative to the observed data, and the dollar amount of the adjustment we make to the outside option.

A common thread across all the counterfactuals with adjusted outside options is that overall welfare is generally reduced relative to the baseline counterfactuals. The reason for this is that the welfare gains achieved in the baseline case almost exclusively result from the shifting of consumers across MA-PD and PDP markets. Many of the baseline counterfactuals result in higher Part D PDP subsidies, which under the adjustments here make the outside option more attractive. A more attractive MA-PD market leads fewer consumers to chose

 $^{^{28}}$ We keep the number of plans fixed to equalize the role of the idiosyncratic error term in the logit model when comparing outcomes across counterfactuals.

²⁹We note that we do not solve for a fixed point where the adjustments in generosity in the Part D PDP market change the MA-PD market, which in turn influences the Part D PDP market, etc. Simulations in a single market suggest that this approximation is very close to the fixed point, as the feedback effects are limited.

the PDP products, which tends to equalize any changes in market share that occurred in the baseline case relative to the observed data. In sum, we end up closer to a situation where the market shares are the same as they are under the observed allocation, but the government is spending much more in subsidy support. Since transfers are socially costly in our framework, this tendency of adjustments to MA-PD to blunt the substitution of MA-PD enrollees to Part D PDP generally lowers overall welfare.

This effect is easily observed in the two counterfactuals where we successively remove links to the LIS and MA-PD markets. With adjustment, the outside option becomes \$44 and \$75 more generous, respectively. In equilibrium, enrollments in Part D PDP are substantially lower, yet per-capita government expenditures in the two programs have increased due to the extra payments to MA-PD under the adjustment to the outside option. This leads to substantially lower overall surplus in both cases.

This story is echoed when considering the effect of market structure on equilibrium outcomes. The case of independent plans looks similar to the counterfactual without LIS or MA-PD links, while monopoly ownership has vastly lower surplus than in the baseline. This is driven by the fact that the monopolist was able to raise the subsidy very high in the baseline counterfactual by raising bids uniformly across all plans. In the adjusted case, this means that the outside option is then made \$111 more attractive. In equilibrium, this leads to a large proportion of consumers switching to MA-PD; payments to MA-PD in turn increase by \$2.5 billion and drive down overall surplus.

The set of counterfactuals examining proportional subsidies also share this feature, although in the case of consumers paying 95 percent of the bid the mechanics work in favor of social surplus. Under the 95 percent counterfactual, consumers pay much more for both PDP and MA-PD plans, while the government pays much less relative to the spending under the observed allocation. This counterfactual highlights in a different way the central story of the paper, which is that the Part D PDP market generates very little willingness-to-pay for the products offered, but primarily increases social welfare through the substitution of consumers from MA-PD. Here, transfers to both markets are minimized and this results in a large increase in social surplus. The same story explains why surplus is much higher with adjustment to the outside option when the government offers a zero voucher. Figure 6 plots total welfare across the original range of vouchers with adjustment to the value of MA-PD market.

6.4 Inclusion of LIS Enrollees

Our baseline set of counterfactuals removes the LIS enrollees from the market. As discussed above, this was done for both pedagogical and technical reasons. Naturally, however, thinking about optimal ways to provide subsidies to both markets is an important question to consider. To shed some light on this issue, we have explored one way of joint subsidy setting for LIS and regular beneficiaries. When thinking about counterfactuals involving LIS enrollees, one must confront two questions. First, what was the role of the LIS random assignment mechanism in determining market outcomes? Second, in the range of counterfactual mechanisms, which one would maximize total surplus across both the regular and LIS markets?

On the first dimension, LIS eligibility threshold plays an important role in disciplining prices in this market. If one does not want LIS enrollees to face premiums or cost-sharing, as is the current situation, then some kind of additional brake on premiums is needed when combining the LIS and regular markets. Without any additional brake, the LIS market by itself effectively functions as a market with 100 percent proportional subsidy, which would lead to unbounded increases in premiums. In the current system, there are two aspects that keep bids down. First, tying LIS and regular markets by imposing that insurers have to submit one bid for both markets emphasizes the elasticity of the regular market. The LIS eligibility threshold serves as another brake on premiums. It introduces an elasticity of demand of sorts through the discontinuity of not being assigned LIS enrollees if the plan's premium is too high, even though the LIS enrollees themselves have zero elasticity of demand for plans that are eligible for LIS random assignment.

In our next set of counterfactuals, we reintroduce LIS enrollees into the regular market and focus on what we consider the most policy-relevant environment: fixed vouchers that can differ across the two market segments. We keep the bid-tying aspect, imposing that insurers set one bid for both markets, but we allow the government to set different subsidies for LIS and regular consumers. We ran the analysis on a matrix of LIS and regular vouchers that range from \$0 to \$1000. Figure 9 illustrates the results. We find that under the double voucher counterfactuals, the optimal voucher for regulars is close to the optimal voucher when just considering the regular market; for the LIS consumers, lower vouchers result in higher surplus due to the high cost of providing the subsidy. The differences across total surplus along the LIS voucher dimension are relatively small as long as the LIS voucher is lower than about \$800, so that setting the LIS voucher to be equal to the optimal voucher for regulars does not lead to a big decrease in total surplus. Note that in this analysis we do not allow for higher social welfare weights on the LIS market; naturally, LIS consumer surplus increases with higher LIS voucher, and hence the optimal set of vouchers would be different under social welfare functions that have higher weights on LIS consumer surplus. The broad takeaway here is that the optimal voucher subsidy policy with and without LIS enrollees is essentially identical.

6.5 Robustness Checks

6.5.1 Vintage and Inertia

Our model of demand includes vintage of the plan as a characteristic in consumer's utility functions. As discussed previously, the interpretation of the coefficient on vintage as a structural parameter will only hold under restrictive circumstances. We view vintage, however, as potentially capturing some inertia in the consumer market. If firms are aware of this inertia, the static Bertrand-Nash pricing game that we assume that they play will result in biased estimates of marginal costs. Recognizing this bias, we perform several robustness checks. Following Miller (2014), who finds that ignoring dynamic incentives can bias marginal costs by a factor of 20 percent, we repeat the key voucher counterfactuals, shifting marginal costs up and down 20 percent relative to our baseline estimates. The results of this exercise are illustrated in Figure 6. The primary result is that the optimal voucher decreases to \$400 when marginal costs are lower, and increases to \$800 when marginal costs are higher. When marginal cost are lower, equilibrium prices are lower, which draws more consumers to the PDP market. Combined with the fact that willingness-to-pay has not changed, the optimal level of subsidy declines. As expected, the case with lower marginal costs results in higher peak welfare, with the situation reversed when marginal costs are 20 percent higher. Consequently, while the levels of welfare estimates could change under a model with dynamic supply-side incentives, the economic intuition driving outcomes across vouchers of varying generosity does not.

Table 10 also reports the effect of changing our estimated marginal costs on welfare under the observed allocation and the social planner. When marginal cost is twenty percent lower, observed surplus increases to \$3.5 billion. On the other side, welfare decreases to negative \$420 million under higher costs. A similar effect occurs in the social planner's case, with the upper bound growing to \$7.9 billion and the lower bound shrinking to \$1.7 billion with adjusted marginal costs. We interpret these two sets of numbers as robust bounds to the amount of surplus generated or possible in this market, respectively.

We also recompute the level of observed and social planner surplus when we exclude the

vintage variable from the welfare calculation. Without the contribution of vintage to utility, observed surplus declines to \$193 million, while the social planner can obtain \$1.2 billion. The inclusion of vintage in the utility function only enters consumer surplus, so in all other counterfactuals that we consider it would result in similar downward shifts in the overall utility.

6.5.2 The Social Cost of Funds

A key parameter in our calculations in the social cost of government funds. In Table 11, we report the overall welfare estimates for each counterfactual in the paper for $\lambda \in \{1, 1.7, 2\}$, in addition to the value of $\lambda = 1.3$ used in our baseline calculations. Increasing the cost of public funds has the general effect of decreasing overall welfare in most cases, although there are a few exceptions, such as the optimal regional voucher and the social planner. The optimal voucher shifts by \$100 at $\lambda = 0$ and down by \$100 at $\lambda = 200$, but otherwise the qualitative ranking of vouchers retains the same general pattern as seen when $\lambda = 1.3$.

We also report the welfare level of the social planner where λ is not applied to the profits of the firm. The idea in our baseline calculations is that the social planner can tax the profits of the firms and use those funds to offset distortionary taxes elsewhere in the economy. Alternatively, one can think of this as the social planner subsidizing losses by firms to obtain gains in consumer surplus, and those subsidy funds again have a social cost. However, this treatment of the social planner problem generates a slight dissonance with the treatment of the overall welfare function, which does not include the λ on profits. Therefore, we recomputed the social planner's problem without the λ on profits, ann report the results of both calculations as a row of Table 11. It turns out that the inclusion or exclusion of the λ on profits makes virtually no difference in the overall welfare calculation, although it does somewhat alter the (unreported) allocation: with costless subsidization of insurers' profits, premiums are slightly lower and hence consumer surplus and inside option enrollment higher. but also insurers' "losses" higher. The reason for the small difference in overall surplus is that the social planner imposes large losses on the firms in both scenarios. The losses are higher when the social planner does not account for the cost of public funds when inflicting losses of firms—this increase counteracts the removal of λ from the profit term in the total surplus calculation.

7 Conclusion

In this paper we have analyzed the welfare effects of the mechanism for determining subsidies for Prescription Drug Plans in Medicare Part D, focusing in particular on the supply side of the market. We draw several conclusions. First, we find that the current program is efficient only if we take into account that the government would have had to subsidize enrollees elsewhere as well. Without taking the latter into account, we could conclude that the program only generates a fraction of dollar value that is spent on it from the federal budget. This is due to a number of related factors: PDP plans are highly subsidized; demand for those plans is generated almost exclusively by said subsidies—consumers have very low willingness-to-pay for unsubsized plans, driven by the availability of inexpensive close substitutes; and firms are able to capture some of the rents of the subsidy mechanism.

On the supply-side we find, perhaps surprisingly, that the current structure of the program, where prices for distinct parts of the program, such as Medicare Advantage Prescription Drug coverage, Low Income Subsidies, and market premiums for regular beneficiaries, are tied together into one mechanism, in fact mutes insurers' ability to raise subsidies, and hence positively affects total welfare. In fact, the current mechanism that weights multiple parts of the program into an average that is used to calculate subsidies, is similar in its incentives to a pre-determined optimal voucher mechanism. We find that providing flat vouchers that are optimally set *ex ante* could increase the total surplus in levels and relative to federal dollars spent, but not by a large amount (although a flat voucher mechanism could dramatically reduce the cost of administering the program, an effect that we do not include in our calculations). We further find that removing the averaging and just setting proportional subsidies would lead to a rapid upward price spiral, as the competitive pressure on the market—which we do find to have reasonably low margins—is, however, not strong enough to mitigate the "raising-the-subsidy" incentives.

Further, our analysis reveals a close connection between Part D and Medicare Advantage that, although not emphasized in the previous literature, proved to be crucial for our findings. We believe that our approach to the quantification of welfare that gradually removes interlinked parts of the environment—specifically, LIS bidding incentives and MA-PD part of the bid average—can be useful for many other public programs that do not exist in isolation, but, instead, are linked to other programs through the choices of consumers and producers or though government transfers.

While our institutional setting focused on the Medicare Part D program, our findings have broader implications for market design of privately-provided and publicly-subsidized social insurance programs. The motivation of subsidizing these programs is typically redistributionthe government attempts to ensure the affordability of insurance. Inevitably, such subsidy policies will have efficiency costs for the market. One source of such inefficiencies is market power. Subsidies create incentives for imperfectly competitive insurers to raise markups and pass them through to the price inelastic government. In general, we show in this paper that conditional on the decision to subsidize social insurance programs, there are large welfare differences across specific mechanisms that are feasibly at the policy maker's disposal. Depending on whether the policy is guided by the considerations of consumer surplus, total welfare, or government spending, different policies deliver drastically different results across these three measures of surplus. Overall, we argue that contrary to the focus of the literature on consumer choices in social insurance markets, the much less studied supply-side behavior in the presence of regulatory intervention and subsidization plays the key role in determining the efficiency outcomes of social insurance programs.

References

- Abaluck, J. and J. Gruber (2011). Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. American Economic Review 101(4), 1180–1210.
- Abaluck, J. and J. Gruber (2016). Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance. American Economic Review 106(8), 2145–2184.
- Abaluck, J., J. Gruber, and A. Swanson (2015). Prescription Drug Use under Medicare PartD: A Linear Model of Nonlinear Budget Sets. NBER Working Paper No. 20976.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile Prices in Market Equilibrium. Econometrica 63(4), 841–889.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. RAND Journal of Economics 25(2), 242–262.
- Center for Medicare and Medicaid Services (2009, February). Guidance to States on the Low-Income Subsidy. http://www.cms.gov/Medicare/Eligibility-and-Enrollment/ LowIncSubMedicarePresCov/EligibilityforLowIncomeSubsidy.html. Accessed June 4, 2015.

- Chorniy, A., D. Miller, and T. Tang (2014). Mergers in Medicare Part D: Decomposing Market Power, Cost Efficiencies, and Bargaining Power. *Princeton University Unpublished Mimeo*.
- Congressional Budget Office (2014, July). Competition and the cost of Medicare's Prescription Drug Program.
- Currie, J. and F. Gahvari (2008). Transfers in Cash and In-Kind: Theory Meets the Data. Journal of Economic Literature 46(2), 333–383.
- Curto, V., L. Einav, J. Levin, and J. Bhattacharya (2015). Can Health Insurance Competition Work? Evidence from Medicare Advantage. NBER Working Paper No. 20818.
- Cutler, D. M. and S. J. Reber (1998). Paying For Health Insurance: The Trade-Off Between Competition And Adverse Selection. *Quarterly Journal of Economics* 113(2), 433–466.
- Decarolis, F. (2015). Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design? American Economic Review 105(4), 1547–1580.
- Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2009). Do switching costs make markets less competitive? *Journal of Marketing Research* 46(4), 435–445.
- Duggan, M. (2004). Does Contracting Out Increase the Efficiency of Government Programs? Evidence from Medicaid HMOs. Journal of Public Economics 88(12), 2549–2572.
- Duggan, M., P. Healy, and F. Scott Morton (2008). Providing Prescription Drug Coverage to the Elderly: America's Experiment with Medicare Part D. Journal of Economic Perspectives 22(4), 69–92.
- Duggan, M. and F. Scott Morton (2006). The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing. Quarterly Journal of Economics 121(1), 1–30.
- Duggan, M. and F. Scott Morton (2010). The Effect of Medicare Part D on Pharmaceutical Prices and Utilization. American Economic Review 100(1), 590–607.
- Einav, L., A. Finkelstein, and M. Polyakova (2016). Private provision of social insurance: drug-specific price elasticities and cost sharing in Medicare Part D. NBER Working Paper No. 22277.

- Einav, L., A. Finkelstein, and P. Schrimpf (2015). The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D. Quarterly Journal of Economics 130(2), 841–899.
- Engelhardt, G. V. and J. Gruber (2011). Medicare Part D and the Financial Protection of the Elderly. *American Economic Journal: Economic Policy* 3(4), 77–102.
- Enthoven, A. C. (2011). Reforming Medicare by Reforming Incentives. *New England Journal* of Medicine 364(21).
- Ericson, K. (2014). Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange. American Economic Journal: Economic Policy 6(1), 38–64.
- Frakt, A. (2011, March). The key difference between exchanges and the Ryan-Rivlin plan. http://theincidentaleconomist.com/wordpress/ the-key-difference-between-exchanges-and-the-ryan-rivlin-plan/.
- Gowrisankaran, G., C. Marsh, and R. Town (2015). Myopia and Complex Dynamic Incentives: Evidence from Medicare Part D. NBER Working Paper No. 21104.
- Gruber, J. and E. Washington (2005). Subsidies to Employee Health Insurance Premiums and the Health Insurance Market. *Journal of Health Economics* 24, 253–276.
- Hausman, J. A. and J. M. Poterba (1987). Household Behavior and the Tax Reform Act of 1986. The Journal of Economic Perspectives 1(1), 101–119.
- Heiss, F., A. Leive, D. McFadden, and J. Winter (2013). Plan Selection in Medicare Part D: Evidence from Administrative Data. *Journal of Health Economics* 32, 1325–1344.
- Heiss, F., D. Mcfadden, and J. Winter (2010). Mind the Gap ! Consumer Perceptions and Choices of Medicare Part D Prescription. In D. Wise (Ed.), *Research Findings in the Economics of Aging*, pp. 413–481. The University of Chicago Press.
- Heiss, F., D. McFadden, J. Winter, A. Wuppermann, and B. Zhou (2016). Inattention and Switching Costs as Sources of Inertia in Medicare Part D. NBER Working Paper No. 22765.
- Heiss, F. and V. Winschel (2008). Likelihood approximation by numerical integration on sparse grids. Journal of Econometrics 144(1), 62-80.

- Ho, K., J. Hogan, and F. Scott Morton (2015). The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program. *NBER Working Paper No. 21028*.
- Hortacsu, A., S. A. Madanizadeh, and S. L. Puller (2015, June). Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market. NBER Working Paper No. 20988.
- Hsu, J., V. Fung, J. Huang, M. Price, R. Brand, R. Hui, B. Fireman, W. H. Dow, J. Bertko, and J. P. Newhouse (2010). Fixing flaws in Medicare drug coverage that prompt insurers to avoid low-income patients. *Health Affairs* 29(12), 2335–2343.
- Kesternich, I., F. Heiss, D. McFadden, and J. Winter (2013). Suit the action to the word, the word to the action: Hypothetical choices and real decisions in Medicare Part D. *Journal* of Health Economics 32(6), 1313–1324.
- Ketcham, J., C. Lucarelli, and C. Powers (2015). Paying Attention or Paying Too Much in Medicare Part D. American Economic Review 105(1), 204–233.
- Ketcham, J. and K. Simon (2008). Medicare Part D's Effects on Elderly Drug Costs and Utilization. American Journal of Managed Care 14(11), 14–22.
- Ketcham, J. D., C. Lucarelli, E. J. Miravete, and M. C. Roebuck (2012). Sinking, Swimming, or Learning to Swim in Medicare Part D. American Economic Review 102(6), 2639–2673.
- Klemperer, P. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *Review* of Economic Studies 62(4), 515–539.
- Kling, J. R., S. Mullainathan, E. Shafir, L. C. Vermeulen, and M. V. Wrobel (2012). Comparison Friction: Experimental Evidence from Medicare Drug Plans. *Quarterly Journal* of Economics 127(1), 199–235.
- Laffont, J.-J. and J. Tirole (1993). A Theory of Incentives in Procurement and Regulation. The MIT Press.
- Lucarelli, C., J. Prince, and K. Simon (2012). The Welfare Impact of Reducing Choice in Medicare Part D: A Comparison of Two Regulation Strategies. *International Economic Review* 53(4), 1155–1177.

- Miller, D. and J. Yeo (2015). The Consequences of a Public Health Insurance Option: Evidence From Medicare Part D. *Clemson University and Singapore Management University Unpublished Mimeo*.
- Miller, D. P. (2016). Risk Adjustment and Low Income Subsidy Distortions in Medicare Part D. *Clemson University Unpublished Mimeo*.
- Miller, D. P. and J. Yeo (2014). A Dynamic Discrete Choice Model of Switching Costs in Medicare Part D. Clemson University and Singapore Management University Unpublished Mimeo.
- Miller, K. (2014). Do Private Medicare Firms Have Lower Costs? University of Oregon Unpublished Mimeo.
- Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Economet*rica 69, 307–342.
- Oliver, T. R., P. R. Lee, and H. L. Lipton (2004). A Political History of Medicare and Prescription Drug Coverage. *Milbank Quarterly* 82(2), 283–354.
- Polyakova, M. (2016). Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D. American Economic Journal: Applied Economics 8(3), 165–95.
- Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica* 49(1), 105–130.
- Starc, A. and R. J. Town (2015). Internalizing Behavioral Externalities: Benefit Integration in Health Insurance. NBER Working Paper No. 21783.
- Town, R. and S. Liu (2003). The welfare impact of Medicare HMOs. *RAND Journal of Economics* 34(4), 719–736.
- Vetter, S., F. Heiss, D. McFadden, and J. Winter (2013). Risk attitudes and Medicare Part D enrollment decisions. *Economics Letters* 119(2), 128–132.
- Williams, H. (1977). On the formation of travel demand models and economic evaluation measure of user benefit. *Environment and Planning* 9(3), 285–344.

- Winter, J., R. Balza, F. Caro, F. Heiss, B. Jun, R. Matzkin, and D. McFadden (2006). Medicare prescription drug coverage: Consumer information and preferences. *PNAS* 103(20), 7929–7934.
- Wu, Y. (2016). Supply Response to Consumer Inertia: Strategic Pricing in Medicare Part D. MIT Unpublished Mimeo.
- Yin, W. and D. Lakdawalla (2015). Insurers' Negotiating Leverage and the External Effect of Medicare Part D. Review of Economics and Statistics 97(2), 314–331.



Notes: Figures illustrate market allocations under proportional (a) and flat or "voucher" subsidy (b) in the stylized duopoly model as described in Section 2. The x-axis in both panels indexes the subsidy value, and is either a fraction of price (in the upper panel) or a normalized monetary amount (in the lower panel). For each subsidy level, we calculate the market share of the outside option (indexed on the left y-axis, as a share); consumer premium (indexed on the left y-axis in normalized monetary units); relative shares between the less and more generous insurance plans (indexed on the left y-axis, as a share); and finally the "excess price" as defined in the text (indexed on the right y-axis in percentage points).

56



Figure 2: Conceptual Model of Subsidies: Welfare Decomposition

(a) Proportional subsidies

Notes: Figures illustrate welfare decomposition under proportional (a) and flat or "voucher" subsidy (b) in the stylized duopoly model as described in Section 2. The x-axis in both panels indexes the subsidy value, and is either a fraction of price (in the upper panel) or a normalized monetary amount (in the lower panel). For each subsidy level, we calculate consumer surplus, produces surplus, and government spending. Using these elements and applying the cost of public funds of 0.3, we calculate total welfare to be the sum of consumer and producer surplus net of government spending.

Figure 3: A stylized structure of a Medicare Part D contract



Notes: Insurers in the Medicare Part D program are required to provide coverage that gives at least the same actuarial value as the Standard Defined Benefit (SDB) that is illustrated in this Figure for years 2007-2010. The SDB design features a deductible, a co-insurance rate of 25 percent up to the initial coverage limit (ICL) and the subsequent "donut hole" that has a 100 percent co-insurance until the individual reaches the catastrophic coverage arm of the contract. The graph illustrates these features of the SDB by mapping the total annual drug spending into the out-of-pocket expenditure. This graph is stylized, since the vast majority of Part D contracts adjust the SDB by either changing cost-sharing within the ICL arm, or adding coverage in the coverage gap.

Figure 4: Contract Prices and Market Structure







Notes: Panel (a) depicts cross-sectional variation in premiums, in the levels of low income subsidies cutoffs (LIPSA), and in "base premium" in 2010. Plans with premiums below LIPSA are eligible for the random assignment of LIS beneficiaries. Panel (b) illustrates the distribution of market shares of the outside and inside options across 34 Medicare Part D regions for regular enrollees.

Figure 5: Estimated Marginal Costs and Mark-ups



Notes: The top panel plots the marginal cost estimates from inversion (for undistorted contracts) and from hedonic projection (for distorted contracts). Plan characteristics used in the hedonic projection include deductible, coverage in the gap and enhanced plan indicators, measures of formulary generosity, pharmacy networks, vintage, as well as estimated unobserved plan quality, and region and insurer fixed effects. The bottom panel shows the mark-ups computed for non-distorted contracts in 2010, using marginal cost estimates from the inversion of insurers' first-order conditions. The mark-up is computed as price net of marginal cost over marginal cost.



Figure 6: Welfare with Uniform Flat Subsidies

Notes: Estimated total welfare in counterfactuals with flat voucher varying in \$100 increments. The vouchers are set to be the same across all geographic markets. In the baseline set of counterfactuals (solid line), optimal uniform voucher is \$600. For comparison, the dashed line plots total surplus for vouchers when the outside option receives the same voucher. We also mark total surplus (and average subsidy) for the observed allocation, the public option counterfactual, as well as the social planner's surplus level. The construction and assumptions behind these counterfactuals are discussed in the text.



Figure 7: Welfare with Uniform Flat Subsidies: Sensitivity to MC estimates

Notes: We adjust the estimated marginal cost by 20% up or down and re-compute flat voucher counterfactuals. This is the margin of error on marginal cost estimates from not accounting for dynamic incentives on the supply side, as found in Miller (2014).

Figure 8: Social Planner's Solution: Changes in Premiums and Market Shares



(a) Premiums

Notes: Differences in premiums (a) and market shares (b) in California market in 2010 between the social planner solution and the observed allocation. On the y-axis, insurance contracts within this market are sorted by increasing estimated social marginal cost.



Figure 9: Welfare with flat vouchers in regular and LIS markets

(a) Graphical

Notes: Total welfare (regular plus LIS markets) in counterfactuals where both regular and LIS enrollees face flat vouchers. We allow for vouchers between the two markets to differ, which generates a matrix of counterfactuals. Panel (a) plots the levels of total welfare relative to the voucher in the regular market for 4 levels of LIS vouchers. Panel (b) reports the matrix of vouchers and their corresponding total welfare in a matrix form.

Table	1:	Summary	Statistics
10010	т.	Summary	0000000000

	Mean (1)	St.Dev. (2)	$ \begin{array}{c} \operatorname{Min} \\ (3) \end{array} $	$\begin{array}{c} \text{Max} \\ (4) \end{array}$
Contracts				
Number of PDP plans	49	5	35	63
Number of large insurers	10	0.2	9	10
Premiums				
Unweighted PDP premium	\$507	\$60	\$375	\$643
Subsidies				
CMS average national bid	\$1.001	\$45	\$965	\$1.060
CMS base consumer premium	\$353	\$26	\$328	\$383
CMS subsidy for average risk beneficiary	\$648	\$20	\$631	\$677
Low income benchmark threshold	\$354	\$24	\$333	\$388
Market-level enrollment, millions				
All Part D Eligible	1.32	1.00	0.06	4.76
PDP enrollment, regular	0.26	0.18	0.01	0.70
PDP enrollment, low-income	0.24	0.20	0.01	1.02
MA-PD enrollment, regular	0.21	0.25	0.00	1.38
MA-PD enrollment, low-income	0.05	0.06	0.00	0.25
Employer sponsored coverage RDS	0.20	0.16	0.01	0.48
Other coverage sources	0.17	0.11	0.01	0.48
No creditable coverage	0.19	0.13	0.01	0.58

Notes: Market-level summary statistics. We have a total of 134 markets (34 geographic regions observed during four years 2007-2010). For each market, we report the average number of plans and large insurers. When counting large insurers we include the following firms separately: Humana, Universal American, WellCare, United Health, CVS Caremark, WellPoint, Coventry, HealthNet, CIGNA. All other insurers are aggregated under "other." National average bids and subsidies do not vary across geographic markets, hence the average is across different years, not geographies. Premiums, subsidies, and bids are all scaled to annual levels.

	Regul	LIS Enr	LIS Enrollees		
	OLS Logit	IV Logit	BLP	OLS Logit	IV Logit
	(1)	(2)	(3)	(4)	(5)
Premium, \$000	-2.74	-10.44		-3.75	-7.59
	(0.15)	(1.09)		(0.12)	(0.47)
Mean, Premium, \$000 (α)			2.58		
			(0.36)		
Std. Deviation, Premium (σ_{α})			0.27		
			(0.09)		
Std. Deviation, Inner Options (σ_{inner})			0.00		
			(0.08)		
Annual Deductible, \$000	-3.25	-6.72	-7.11		
	(0.21)	(0.56)	(1.52)		
Gap Coverage	0.18	2.93	2.89		
	(0.07)	(0.40)	(0.94)		
Std. Deviation, Gap Coverage (σ_{qap})	. ,	. ,	0.00		
			(0.08)		
Number of Top Drugs Covered	0.24	31.58	30.4	-1.42	11.38
	(4.36)	(7.30)	(12.1)	(5.49)	(7.36)
Pharmacy Network Measure	0.31	0.29	0.31	0.22	0.23
	(0.04)	(0.06)	(0.01)	(0.06)	(0.09)
Number Years Plan on Market	0.61	0.88	0.90	0.78	0.83
	(0.02)	(0.05)	(0.12)	(0.29)	(0.04)

Table 2: Demand Estimates

Notes: Standard errors are reported in parentheses. The sample is all Medicare Part D stand-alone prescription drug plans from 2007 to 2010; the sample size is 6,675 plan-year observations for regular enrollees and 4,860 planyear observations for LIS enrollees. The LIS sample size is smaller, because we aggregate all plans eligible for LIS random assignment into one plan. In addition to the displayed coefficients and fixed effects, all models also include a constant and the following plan characteristics: a dummy for an enhanced plan, the number of APIs in the plan's formulary, and the number of drugs placed in Tiers 1-2 of the formulary which have low cost-sharing. Columns (2), (3), and (5) use the set of instrumental variables described in the text. The first stage for these instruments is reported in Table 3; the F-statistic for the first stage in the regular market is 247, in the LIS market it is 23.

	$\begin{array}{c} \text{Regular} \\ (1) \end{array}$	$\begin{array}{c} \text{LIS} \\ (2) \end{array}$
No. of PDP plans in a region-year by same PO	$\begin{array}{c} -9.631^{***} \\ (2.375) \end{array}$	2.048 (4.101)
No. of MA plans in a region-year by same PO	-0.184 (0.242)	-0.467 (0.470)
Deductible of MA plans in the same region-year	-0.147 (0.173)	-0.0204 (0.326)
Hausman IV	$\begin{array}{c} 0.371^{***} \\ (0.0326) \end{array}$	$\begin{array}{c} 0.721^{***} \\ (0.0471) \end{array}$
Observations	6675	5313

Table 3: First Stage Regression for Plan Premiums

Standard errors in parentheses

* p < 0.05, ** p < 0.01, *** p < 0.001

Notes: Standard errors in reported in parentheses. Each regression uses data on Medicare Part D stand-alone prescription drug plans in years 2007 to 2010. See the text for more details on the construction of the instruments. The regressions also include all plan characteristics that are used in demand estimation, including a constant and fixed effects for geographic markets, parent organizations, and years.

	Berry MC inversion (1)	BLP MC inversion (2)
Annual deductible	-0.365^{***} (0.0440)	-0.354^{***} (0.0428)
No. of common APIs	$\begin{array}{c} 0.142 \\ (0.130) \end{array}$	$0.183 \\ (0.127)$
Has coverage in the gap $(1/0)$	$\begin{array}{c} 0.422^{***} \\ (0.0105) \end{array}$	$\begin{array}{c} 0.412^{***} \\ (0.0102) \end{array}$
Enhanced plan $(1/0)$	-0.0352^{**} (0.0118)	-0.0326^{**} (0.0114)
No. of top drugs in Tier 1 and 2	-0.569 (0.380)	-0.422 (0.370)
No. of top drugs covered	-8.696^{***} (2.476)	-8.975^{***} (2.407)
Pharmacy network measure	-0.188^{***} (0.0509)	-0.196^{***} (0.0495)
Number of years the plan is on the market	$ \begin{array}{c} 46.09^{***} \\ (3.170) \end{array} $	$44.01^{***} \\ (3.082)$
Mean dep.var.	1.171	1.172
Std. dev. dep. var.	0.239	0.228
R-squared	0.868	0.864
Number of observations	756	756

Table 4: Marginal Cost Projection

Standard errors in parentheses

* p < 0.05, ** p < 0.01, *** p < 0.001

Notes: Main coefficients for the hedonic projection of estimated marginal costs on the characteristics of plans for "non-distorted" insurance plans as defined in the text. We restrict the model to year 2010, as that is the year in which we do the counterfactual analyses. For BLP model, we use the Berry IV specification for the LIS part of the market when inverting the first order conditions. Fixed effects for markets and insurers are included in the regression but not reported.

	Observed Allocation		No LIS	No LIS or	Independent	Monopoly
	Regular (1)	LIS (2)	$ \begin{array}{c} \operatorname{Link}\\ (3) \end{array} $	MA-PD Link (4)	$\begin{array}{c} \text{Plans} \\ (5) \end{array}$	Ownership (6)
Consumer Surplus, \$M	2,517	2,809	2,881	3,399	3,336	2,807
Insurer Profit, \$M	447	-	1,058	1,275	1,168	2,151
Premium Subsidy, \$M	5,936	8,067	7,768	10,082	$9,\!695$	8,853
Reinsurance Subsidy, \$M	$3,\!444$	2,588	4,042	4,960	4,924	4,049
Inside Option Enrollment, '000	8,772	7,829	10,512	12,892	$12,\!580$	10,556
Inside Option, Percent	38	75	45	55	54	45
Average Weighted Premium, \$	502	21	473	420	428	450
Average Weighted Bid, \$	1,123	1,051	$1,\!153$	$1,\!145$	1,142	1,241
MA-PD Premium Subsidy, \$M	6,018	$5,\!371$	7,211	8,844	$8,\!630$	7,241
MA-PD Reinsurance Subsidy, \$M	2,281	5,566	2,733	3,352	3,271	2,745
Total surplus, \$M	$1,\!559$	3,175 ^a	1,515	974	970	$1,\!167$
per PDP Subsidy Dollar, \$	0.17	0.30	0.13	0.06	0.07	0.09
per Opportunity Cost Dollar, \$	1.11	-	0.62	0.26	0.27	0.31

Table 5: Observed Allocation; Link and Ownership Counterfactuals

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures if all PDP enrollees switched to MA-PD. Welfare per opportunity cost dollar is only reported when government expenditures in PDP are higher than in MA-PD.

^aCalculation assumes zero insurer profit.

	Proportional Subsidies			Vouchers		
	p=5% (1)	p=32% (2)	p=95% (3)	\$0 (4)	\$700 (5)	\$1,400 (6)
Consumer Surplus, \$M	13,344	4,533	1,081	1,073	2,497	12,581
Insurer Profit, \$M	$50,\!405$	4,845	20	13	879	4,814
Premium Subsidy, \$M	$84,\!156$	18,086	10	-	$5,\!997$	32,599
Reinsurance Subsidy, \$M	$15,\!695$	8,183	58	39	$3,\!296$	10,367
Inside Option Enrollment, '000	23,280	14,498	133	89	8,567	23,297
Inside Option, Percent	100	62	1	0.4	37	100
Average Weighted Premium, \$	190	587	1,382	$1,\!441$	522	77
Average Weighted Bid, \$	$3,\!805$	1,835	$1,\!455$	$1,\!441$	1,222	1,477
MA-PD Premium Subsidy, \$M	15,970	9,946	92	61	$5,\!877$	$15,\!982$
MA-PD Reinsurance Subsidy, \$M	6,053	3,770	35	23	2,227	6,057
Total surplus, \$M	$(37,\!428)$	(6, 942)	$1,\!177$	$1,\!146$	1,830	(9,811)
per PDP Subsidy Dollar, \$	(0.37)	(0.26)	17.29	29.62	0.20	(0.23)
per Opportunity Cost Dollar, \$	(0.37)	(0.43)	-	-	1.18	(0.36)

Table 6: Proportional and Flat Subsidy Counterfactuals

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures that would be incurred if all PDP enrollees switched to MA-PD. Welfare per opportunity cost dollar is only reported when government expenditures in PDP are higher than in MA-PD.
	Margina	al Cost	Soci	al Planner	Optimal Voucher		
	Private (1)	Social (2)	Baseline (3)	Public Option (4)	National (5)	Market (6)	
Consumer Surplus, \$M	1,094	1,064	2,871	2,996	1,819	1,934	
Insurer Profit, \$M	-	-	(9,725)	-	514	561	
Premium Subsidy, \$M	-	-	-	8,785	2,852	3,332	
Reinsurance Subsidy, \$M	93	-	-	$3,\!655$	1,842	1,948	
Inside Option Enrollment, '000	216	35	10,895	12,996	4,753	$5,\!272$	
Inside Option, Percent	1	0	47	56	20	23	
Average Weighted Premium, \$	1,284	$1,\!450$	377	192	650	631	
Average Weighted Bid, \$	1,284	$1,\!450$	377	868	1,250	1,246	
MA-PD Premium Subsidy, \$M	148	24	$7,\!474$	8,915	3,260	$3,\!617$	
MA-PD Reinsurance Subsidy, \$M	56	9	2,833	$3,\!379$	$1,\!236$	$1,\!371$	
Total surplus, \$M	$1,\!240$	$1,\!106$	$3,\!627$	$2,\!806$	$2,\!075$	$2,\!116$	
per PDP Subsidy Dollar, \$	13.40	-	0.29	0.23	0.44	0.40	
per Opportunity Cost Dollar, \$	-	-	-	-	8.08	5.58	

Table 7: Marginal Cost, Social Planner, and Optimal Voucher Counterfactuals

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures that would be incurred if all PDP enrollees switched to MA-PD. Welfare per opportunity cost dollar is only reported when government expenditures in PDP are higher than in MA-PD. The public option counterfactual in Column (4) assumes that in the inside option, individuals are only offered a plan that has marginal cost equal to the lowest marginal cost PDP plan that is observed in the data (separately for each market); the premiums for these plans are equal to marginal cost net of observed subsidy level.

	Data (1)	No LIS Link (2)	No LIS or MA-PD Link (3)	Independent Plans (4)	Monopoly Ownership (5)
Consumer Surplus, \$M	2,517	$3,\!483$	4,309	4,151	4,709
Insurer Profit, \$M	447	856	911	867	1,199
Premium Subsidy, \$M	5,936	6,152	6,970	6,927	$5,\!651$
Reinsurance Subsidy, \$M	$3,\!444$	3,204	3,429	3,506	2,585
Inside Option Enrollment, '000	8,772	8,325	8,913	8,988	6,738
Inside Option, Percent	38	36	38	38	29
Average Weighted Premium, \$	502	484	438	445	444
Average Weighted Bid, \$	$1,\!123$	1,162	$1,\!159$	$1,\!156$	1,229
MA-PD Premium Subsidy, \$M	6,018	5,711	$6,\!115$	6,166	$4,\!622$
MA-PD Reinsurance Subsidy, \$M	2,281	2,165	2,317	2,337	1,752
Additional MA-PD Subsidy Payment, \$M	-	797	1,387	1,222	2,544
MA-PD Change in Average Premium, $\$$	-	(44)	(75)	(67)	(111)
Total surplus, \$M	$1,\!559$	$1,\!379$	859	921	181

Table 8: Observed Allocation; Link and Ownership Counterfactuals with Endogenous Outside Option

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures that would be incurred if all PDP enrollees switched to MA-PD. MA-PD additional payment reflects the amount of additional subsidies that the government has to pay when MA-PD subsidies are adjusted to the level of PDP subsidies. The MA-PD change in average premium reflects the average change in consumer premiums for the outside option when the outside option is adjusted to receive the same subsidies as the PDP market. For all counterfactuals in this table, the subsidy adjustment for the outside option is at the subsidy level, i.e. we calculate the average counterfactual subsidy in PDP market and apply it to MA-PD bids.

	Propo	rtional Sub	sidies			
	p=5% (1)	p=32% (2)	p=95% (3)		\$700 (5)	\$1,400 (6)
Consumer Surplus, \$M	$13,\!389$	4,358	(12,304)	(13, 496)	2,720	12,704
Insurer Profit, \$M	49,008	5,129	851	778	815	4,414
Premium Subsidy, \$M	$81,\!651$	19,262	486	-	5,519	30,250
Reinsurance Subsidy, \$M	15,161	8,735	3,030	2,886	3,035	$9,\!629$
Inside Option Enrollment, '000	$22,\!488$	15,518	7,790	7,492	7,884	$21,\!620$
Inside Option, Percent	96	66	33	32	34	93
Average Weighted Premium, \$	191	584	1,184	1,228	526	80
Average Weighted Bid, \$	3,822	1,825	1,247	1,228	1,226	1,479
MA-PD Premium Subsidy, \$M	15,427	$10,\!645$	5,344	5,140	5,408	14,831
MA-PD Reinsurance Subsidy, \$M	5,847	4,035	2,025	1,948	2,050	$5,\!621$
Additional MA-PD Subsidy Payment, \$M	202	(235)	(9,932)	(10,887)	217	1,244
MA-PD Change in Average Premium, \$	(152)	22	628	677	(-15)	(-222)
Total surplus, \$M	(36,064)	$(7,\!519)$	$6,\!468$	$6,\!897$	1,828	(9,753)

Table 9: Proportional and Flat Subsidy Counterfactuals with Endogenous Outside Option

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures that would be incurred if all PDP enrollees switched to MA-PD. MA-PD additional payment reflects the amount of additional subsidies that the government has to pay when MA-PD subsidies are adjusted. The MA-PD change in average premium reflects the average change in consumer premiums for the outside option when the outside option is adjusted. In this table, the subsidy adjustment for the outside option is done by applying the same subsidy rules rather than subsidy levels across markets. In other words, we apply a 5%, 78%, and 95% subsidy to MA-PD bids rather than equalizing subsidy levels between PDP and MA-PD. For the voucher case, we apply the same voucher to both PDP and MA-PD markets.

	Observed Allocation			Social Planner			
	No vintage (1)	MC -20% (2)	MC + 20% (3)	No vintage (4)	MC -20% (5)	MC + 20% (6)	
Consumer Surplus, \$M	1,152	2,517	2,517	1,198	$5,\!114$	1,572	
Insurer Profit, \$M	447	2,426	(1,531)	(1,221)	(15,908)	(3,189)	
Premium Subsidy, \$M	$5,\!936$	5,936	5,936	_	_	-	
Reinsurance Subsidy, \$M	$3,\!444$	$3,\!444$	$3,\!444$	-	-	-	
Inside Option Enrollment, '000	8,772	8,772	8,772	1,325	$19,\!145$	3,508	
Inside Option, Percent	38	38	38	6	82	15	
Average Weighted Premium, \$	502	502	502	294	244	563	
Average Weighted Bid, \$	$1,\!123$	1,123	1,123	294	244	563	
MA-PD Premium Subsidy, \$M	6,018	6,018	6,018	909	$13,\!134$	2,407	
MA-PD Reinsurance Subsidy, \$M	2,281	2,281	2,281	345	4,978	912	
Total surplus, \$M	193	$3,\!537$	(420)	1,240	7,978	1,740	

Table 10: Sensitivity checks

Notes: Total surplus is the sum of consumer surplus and producer profit minus the social opportunity cost of government spending. MA-PD premium and reinsurance subsidies are the amount of government expenditures that would be incurred if all PDP enrollees switched to MA-PD. In columns (1) and (4) we remove vintage from the utility function when computing welfare, but not when determining the allocations. In other words, we allow for inertial choices to mute demand responsiveness, but treat this measure of inertia as welfare-neutral. In columns (2)-(3) and (5)-(6) we adjust the estimated marginal cost by 20% up or down. This is the margin of error on marginal cost estimates from not accounting for dynamic incentives on the supply side, as found in Miller (2014).

	$\lambda = 1$	$\lambda = 1.3$	$\lambda = 1.7$	$\lambda = 2$
	(1)	(2)	(3)	(4)
Observed allocation	1,883	1,559	1,126	802
Observed allocation with welfare-neutral vintage	518	193	(239)	(563)
Observed allocation with MC 20% down	3,861	$3,\!537$	3,104	2,780
Observed allocation with MC 20% up	(96)	(420)	(853)	(1, 177)
No LIS link	2,074	1,515	768	209
No LIS or MA-PD link	1,828	974	(165)	(1,019)
Independent Plans	1,785	970	(118)	(933)
Monopoly Ownership	2,042	1,167	0	(875)
5 percent premium	(14,079)	(37, 428)	(68, 559)	(91, 907)
32 percent premium	(3,176)	(6,942)	(11,964)	(15,730)
95 percent premium	1,159	1,177	1,200	1,217
Vouchers				
\$0	1,132	1,146	1,165	$1,\!178$
\$100	1,182	1,204	1,232	1,254
\$200	1,266	1,298	1,342	1,375
\$300	1,400	1,446	1,509	1,555
\$400	1,600	$1,\!656$	1,731	1,787
\$500	1,867	1,907	1,960	1,999
\$600	$2,\!134$	2,075	$1,\!996$	1,937
\$700	$2,\!187$	$1,\!830$	$1,\!355$	998
\$800	$1,\!685$	686	(645)	(1, 644)
\$900	600	(1,373)	(4,004)	(5,978)
\$1,000	(395)	(3,370)	(7, 337)	(10, 311)
\$1,100	(571)	(4,286)	(9,239)	(12,954)
\$1,200	(654)	(5,037)	(10,882)	(15, 265)
\$1,300	(1, 482)	(6, 695)	(13, 645)	(18,858)
\$1,400	$(3,\!533)$	(9,811)	(18, 182)	(24, 460)
\$1,500	(6, 198)	(13, 578)	(23, 418)	(30,798)
Optimal regional voucher	$2,\!254$	$2,\!116$	2,064	2,067
Marginal Cost	1,206	$1,\!240$	1,285	1,318
Social Marginal Cost	1,097	$1,\!106$	$1,\!119$	1,129
Social Planner				
Baseline	$3,\!452$	$3,\!627$	3,859	4,034
Welfare-neutral vintage	1,230	$1,\!240$	1,253	1,262
MC 20% down	$7,\!317$	$7,\!978$	8,860	$9,\!521$
MC 20% up	1,701	1,740	1,792	$1,\!831$
Public option	$2,\!850$	$2,\!806$	2,748	2,704
No cost of public funds on profits	$3,\!517$	$7,\!306$	$12,\!359$	$16,\!149$

Table 11: Sensitivity to the cost of public funds parameter

Notes: The table reports total surplus level for all baseline counterfactuals at different levels of the cost of public funds. Column (2) with the cost of public funds equal to 0.3 corresponds to the preferred specification reported in earlier tables.

8 Appendix

8.1 Derivation of Social Planner's Pricing Problem

Social planner's problem is to maximize the following welfare function:

$$W^{SP}(p) = \int \frac{1}{\alpha} \left(\gamma + \ln \left[1 + \sum_{j=1}^{J} \exp(v_j(\theta), p_j) \right] \right) dF(\theta) + \lambda \left[\underbrace{\sum_{j=1}^{J} \left(p_j - c_j \right) s_j(p)}_{j=1} - \underbrace{\left(\sum_{j=1}^{J} \left(G_j^{PDP} - G^{MAPD} \right) s_j(p) \right)}_{j=1} \right].$$
(19)

The social planner's solution is defined by the set of first-order conditions obtained by differentiating $W^{SP}(p)$ with respect to prices. The derivative of consumer surplus with respect to p_j has a conveniently simple form after some simplifications:

$$\frac{\partial CS(p)}{\partial p_j} = \int \frac{1}{\alpha} \left[\frac{-\alpha \exp(v_j(\theta))}{1 + \sum_{k=1}^J \exp(v_k(\theta))} \right] dF(\theta) = -s_j(p).$$
(20)

The derivative of product market profit with respect to p_j is:

$$\frac{\partial \Pi(p)}{\partial p_j} = \lambda s_j(p) + \lambda \sum_k (p_k - c_k) \frac{\partial s_k(p)}{\partial p_j}.$$
(21)

The derivative of government spending with respect to p_j is:

$$\frac{\partial GS(p)}{\partial p_j} = -\lambda \left[\sum_k (G_k^{PDP} - G^{MAPD}) \frac{\partial s_k(p)}{\partial p_j} \right],\tag{22}$$

$$= -\lambda \left[\sum_{k} \Delta G_k \frac{\partial s_k(p)}{\partial p_j} \right].$$
(23)

Summing these terms, we obtain:

$$\frac{\partial W(p)}{\partial p_j} = (\lambda - 1)s_j + \lambda \sum_k (p_k - c_k) \frac{\partial s_k(p)}{\partial p_j} - \lambda \left[\sum_k \Delta G_k \frac{\partial s_k(p)}{\partial p_j} \right], \tag{24}$$

$$= (\lambda - 1)s_j + \lambda \sum_k (p_k - c_k - \Delta G_k) \frac{\partial s_k(p)}{\partial p_j}.$$
 (25)

Note that a decrease in consumer surplus in response to an increased price $(-s_j(p))$ is offset, up to the cost of transferring public funds, by an increase in profit in the product market $(\lambda s_j(p))$.

The first-order conditions can be expressed in a particularly simple formula in vector notation; the set of equations defining the social planner's solution is:

$$(\lambda - 1)s(p) + \lambda\Omega(p)(p - c - \Delta G) = 0, \qquad (26)$$

where $\Omega(p)$ is a matrix of partial derivatives such that the element in the *i*-th row and *j*-th column is:

$$\Omega_{ij}(p) = \frac{\partial s_j(p)}{\partial p_i}.$$
(27)

8.2 Vintage and a Two-Stage Model of Inattention and Choice

Our demand system contains a vintage variable to proxy for behavioral biases in demand arising from inattention and switching costs. In this appendix, we show how one can derive the reduced form specification from an explicit structural model. We start by borrowing from Hortacsu et al. (2015), who posit a two-stage model of choice with inattention. In the first stage consumers make an active choice with probability α . In the second stage, attentive consumers face a standard discrete choice problem, while inattentive consumers stay in the same plan that they had in the last period. This implies that the observed share of plan jdepends on its own share from the previous period as follows:

$$\hat{s}_{j,t}(p, s_{j,t-1}) = \alpha M Pr_{j,t}(p) + (1 - \alpha)s_{j,t-1},$$
(28)

where \hat{s} is the observed share, p is the vector of plan premiums, M is market size, and $Pr_{j,t}(p)$ is the usual logit probability. In the first year of the program, this model reduces to the usual logit model, or equivalently, our discrete choice model with vintage set to zero. In year two, the observed share is a convolution of the current choice share and the set of

inattentive consumers who did not make a choice. Irrespective of whether p = 0, where no one pays attention, or p = 1, where everyone is perfectly attentive, the plan accumulates consumers as time goes on and the relative share of the plan remains fixed as the rest of the world stays constant. The distinguishing feature of this model, however, is that this model predicts that the firm can start raising premiums after the first year without losing as much market share as it would have in a perfectly attentive world. To see this, the derivative of Equation 28 with respect to its premium only has the current set of active choosers in it:

$$\frac{\partial \hat{s}_{j,t}(p, s_{j,t-1})}{\partial p_{j,t}} = \alpha M \frac{\partial Pr_{j,t}(p)}{\partial p_{j,t}},\tag{29}$$

while profits are a function of the total share, of which fraction $(1 - \alpha)$ are unresponsive to price changes. The key point is that as α declines, the firm can increasingly raise premiums and retain the same market share.

The mapping from this model to our model with a vintage variable is direct: as the market evolves, the share of active choosers effectively shrinks as an increasing percentage of consumers have been in the market for longer than one period. In the simplest case, assuming that no one exits the market and all pre-existing consumers are completely inattentive, $\alpha(T) = 1/(T-1)$, where T is the number of periods the market has been active. Our vintage variable proxies directly for this effect, as one can rewrite Equation 29 as:

$$\frac{\partial \hat{s}_{j,t}(p, s_{j,t-1})}{\partial p_{j,t}} = \alpha(T)M\frac{\partial Pr_{j,t}(p)}{\partial p_{j,t}} = M\frac{\partial Pr_{j,t}(p, \beta(T))}{\partial p_{j,t}},$$
(30)

where $\beta(T) = F(\alpha(T))$ is a positive, monotonic transformation of $\alpha(T)$. This mapping can be generalized to allow for where $\alpha > 0$ for pre-existing consumers or where α is a function of the premium change (a la Heiss et al. (2016) and Ho et al. (2015)). As such, one can view our reduced form model of demand with a vintage variable as arising from a structural two-stage model of inattention and choice.

8.3 Algorithm for Solving Counterfactual Equilibria

Several of our counterfactuals involve resolving equilibrium bids when the subsidy is an endogenous function of the average bid. We solve these types of equilibria in a nested fixed point algorithm. In the outer step, we first pose an average bid. We model the firms as taking this average bid as fixed. This is not an unreasonable assumption, as the marginal effect of any one firm's bids on the average bid is going to be very small, as the bid is a function of 1500 plans (and possibly MA-PD enrollment-weighted prices). Taking this average bid, b, as fixed, we then solve for the vector of first-order conditions among all plans that we classify as not seeking LIS random assignment enrollees. After finding this vector of bids across all markets, we then compute the enrollment-weighted average bid. We grid search over a range of average bids until we find an average bid that correctly reflects the equilibrium average bid.

We also note that we use the sparse grids method described in Heiss and Winschel (2008) for the evaluation of all integrals. Sparse grids are efficient and accurate multidimensional quadrature methods with excellent performance. Estimation of the BLP specifications was standard with the exception of imposing the lognormality of the price coefficient. All codes are written in Java and are publicly available, as is the data required to estimate all models and compute counterfactuals.

8.4 Calculation of MA-PD adjustments

We proceed in several steps to calculate the adjustments to MA-PD value in the counterfactuals that allow for changes in the outside option. There are slight differences across different counterfactuals, so we describe them separately. We start with counterfactuals reported in Table 8. These counterfactuals compute PDP subsidies via bid-averaging, similar to how the subsidies are calculated under the observed allocation. To compute MA-PD adjustments for these counterfactuals, we start with their unadjusted counterparts in Table 5. In the first step, for each of these counterfactuals, we compute the counterfactual premium subsidy for PDP plans. This subsidy just represents the difference between PDP bids and premiums. Next, we turn to the MA-PD market. In the data, we observe MA-PD premiums. but we do not observe bids. We impute MA-PD bids in the same way we imputed PDP bids, but inverting the subsidization formula. One complication is that MA-PD plans can apply additional subsidies to their bids, by pulling in resources from the medical part of the Medicare Advantage program. MA plans can use their MA subsidy to "buy down" MA-PD premiums. The data on these "buy down" are sparse, and we do not have it at individual plan level. To address this concern, we turn to the MA literature, specifically Kluender and Mast (2016), who report that the average MA-PD buy down is \$3.90 a month. We apply this adjustment to all MA-PD premiums, so that our imputed monthly MA-PD bid is equal to basic MA-PD premium observed in the data plus \$3.90 adjustment plus \$88.33 (which is the national average bid that was released by CMS in 2010) and minus \$31.94 (which was the base beneficiary premium released by CMS in 2010). With these imputed MA-PD bids in hand, we apply the counterfactual PDP subsidy to each of these bids. In addition, we keep the idea that MA-PD plans would apply the same "buy-down" on top of any Part D subsidy. Subtracting the counterfactual subsidy as well as the \$3.90 "buy-down" from each MA-PD bid gives us counterfactual MA-PD premiums for each MA-PD plan. In many cases the counterfactual PDP subsidy together with the "buy-down" are higher than the MA-PD bid. In these cases we impose a zero lower bound on MA-PD premiums — this is in line with the observed allocation in which many MA-PD plans have zero premiums. In the next step, for each MA-PD plan we compute the difference between the observed and counterfactual premium. We take the average of these differences across all plans, which gives us one number that summarizes the average change in MA-PD premiums under the application of the counterfactual PDP subsidy. This number is recorded as a separate row in Table 8 as "MA-PD Change in Average Premium". We use this (annual) number as our measure of the change in the value of the outside option.

Computations in Table 9 are in general very similar. The only difference is in how we compute the counterfactual Part D subsidy. In the case of vouchers, this is simple, since we just apply the same voucher level to MA-PD as we did to PDP. In the case of proportional subsidies, we do not calculate the counterfactual PDP subsidy and instead directly apply the counterfactual subsidization rule to MA-PD bids. For example, in the case of 5% premium, we let counterfactual MA-PD premiums be equal to MA-PD bids multiplied by 5%.

8.5 Alternative Demand Estimates

Coefficient	Point Estimate	SE
Without Any Random Coefficients		
σ_{price}	-	-
σ_{gap}	-	-
$ar{lpha}$	2.32	0.208
σ_{inside}	-	-
GMM Objective Value: 0.00377880		
With RC on Coverage in the Gap		
σ_{price}	0.268	0.111
σ_{gap}	0.00	0.0836
$\bar{\alpha}$	2.58	0.257
σ_{inside} GMM Objective Value: 0.00332521	-	-
With RC on the Inside Option		
σ_{price}	0.268	0.0377
σ_{gap}	-	-
ā	2.58	0.265
σ_{inside}	0.00	0.0460
GMM Objective Value: 0.00332521		
With RC on Coverage in the Gap and Inside Option		
σ_{price}	0.268	0.0742
σ_{gap}	0.00	0.0277
\bar{a}	2.58	0.220
σ_{inside}	0.00	0.0202
GMM Objective Value: 0.00332521		

Table 12: Alternative Demand Estimates for Regular Market