NBER WORKING PAPER SERIES

SUBSIDY DESIGN IN PRIVATELY-PROVIDED SOCIAL INSURANCE: LESSONS FROM 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, Bart Hamilton, Ben Handel, Kate Ho, Thomas Jeitschko, Claudio Lucarelli, Aviv Nevo, Ariel Pakes, Nancy Rose, Marc Rysman, Bernardo Silveira, Chris Timmins, 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) and the National Institute on Aging (5P01AG005842-29). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Subsidy Design in Privately-Provided Social Insurance: Lessons from Medicare Part D
Francesco Decarolis, Maria Polyakova, and Stephen P. Ryan
NBER Working Paper No. 21298
June 2015, Revised September 2018
JEL No. H2,H4,I11,I18,L1,L2

ABSTRACT

The efficiency of publicly-subsidized, privately-provisioned social insurance programs depends on the interaction between strategic insurers and the subsidy mechanism. We study this interaction in the context of Medicare's prescription drug coverage program. We find that the observed mechanism is successful in keeping "raise-the-subsidy" incentives relatively low, acts much like a at voucher, and obtains a level of welfare close to the optimal voucher. Across a range of counterfactuals, we find that more efficient subsidy mechanisms share three features: they retain the marginal elasticity of demand, limit the exercise of market power, and preserve the link between prices and marginal costs.

Francesco Decarolis
Bocconi University
Via Sarfatti 25
Milan, 20100
Italy
and EIEF
francesco.decarolis@unibocconi.it

Maria Polyakova
Department of Health Research & Policy
Stanford University
Redwood Building T111
150 Governor's Lane
Stanford, CA 94305
and NBER
maria.polyakova@stanford.edu

Stephen P. Ryan
Olin Business School
Washington University in St. Louis
Campus Box 1133
One Brookings Drive
St. Louis, MO 63130
and NBER
stephen.p.ryan@wustl.edu

1 Introduction

Social insurance programs have traditionally been provided directly by the government. The last two decades, however, have seen an accelerating effort to move these programs to privately-provided markets with public funding and extensive regulatory oversight (McGuire et al., 2011; Gruber, 2017). Recent examples in the United States include the use of private managed care plans in Medicare and Medicaid, private insurers competing for consumers under the Health Insurance Marketplaces of the Affordable Care Act, and, the subject of this paper, the private provision of prescription drug benefits in Medicare Part D. This trend has also extended around the world to, for example, the privatization of social security benefits in Chile, and a discussion of privatizing social security, disability, and unemployment insurance systems in many OECD countries. The broad objectives of all such programs is to leverage the benefits of competition to provide high-quality services at low cost to both consumers and the government.

One critical issue that arises in such settings is how governments should determine the level and the distribution of public funding that flows into these programs. The funding is typically channeled through subsidies that are paid directly to private firms, making the question of how to tailor the subsidy mechanisms central. While answering this question in general is beyond the scope of any single paper, we contribute to the emerging academic and policy discussion of the subsidy mechanism by providing novel evidence from the Medicare Part D Prescription Drug Plan (PDP) program. Part D is an elective prescription drug insurance program available to Medicare beneficiaries that was launched in 2006 and since then has become one of the role models for privately-provided publicly-financed social insurance programs in the United States. The PDP market has several features which make it well-suited for studying subsidy mechanism design: the program has clear, well-articulated rules that allow us to cleanly model the incentives facing strategic firms; excellent data exists on potential consumers and the set of choices available to them; and a complex mechanism links equilibrium market outcomes to consumer-facing plan prices and public subsidies.

Using this rich environment, we analyze equilibrium allocations, prices, and the incidence of subsidy dollars under the existing subsidization mechanism, as well as under an array of counterfactual subsidy mechanisms that resemble a variety of policy proposals in this and related markets. To facilitate this analysis, we proceed in two steps. We posit and estimate a structural model of consumer demand and strategic insurers, respecting the many institutional details present in the market. On the demand side, we allow for risk-based selection by allowing preferences and costs to vary across six different consumer types. Allowing for risk-based selection enables us to directly model the key difference between insurance markets and regular product markets: the marginal cost curve is an endogenous function of equilibrium prices. On the supply side, we build a profit function for insurers, accounting for a host of details such as the subsidy mechanism, risk-specific payments, reinsurance, multiple demand types, and the endogenous marginal cost function. We leverage this structural model to first estimate demand and cost primitives before turning to counterfactual simulations where we adjust the subsidy mechanism.

We find that equilibrium outcomes, and, by extension, welfare, are driven by four empirical facts. First, we estimate that consumers have relatively low intrinsic willingness-to-pay for independent prescription drug plans. This is primarily driven by the existence of a highly-subsidized close substitute - prescription drug coverage bundled with medical insurance under private Medicare managed care plans known as Medicare Advantage (MA-PD). The second, related, result is that the primary driver of welfare is the opportunity cost of government spending. We find that the sign and magnitude of our welfare estimates are dominated by the ability of the government to set subsidies to achieve optimal sorting of consumers, and risks, across different types of prescription drug coverage. Third, the observed subsidy-setting mechanism appears to be successful in keeping plans' margins relatively low, as insurers price near marginal cost. Fourth, we find that, once one distills all of the administrative details of how the Part D market works, that the current mechanism acts much like a flat voucher and obtains a level of welfare close to the optimal voucher. However, we estimate that a social planner could do substantially better, by adjusting prices to correct our finding that consumers are purchasing too few and too socially-expensive plans relative to the social optimum.

While our estimates are specific to the context of Part D, our results suggest several key economic forces that are likely to be important in any setting with publicly-subsidized privately-provisioned goods and services. First, it is important to preserve the marginal relationship between the prices that firms set and the prices that consumers pay. This keeps the elasticity of demand relatively high, which results in more intense competition and lower prices in equilibrium. Second, consumer-facing prices should be positively related to the social cost of providing those services. Third, the relationship between subsidies and equilibrium outcomes needs to be carefully tempered to prevent strategic pricing by imperfectly competitive firms.

Our paper is related to several distinct literatures in social insurance, design of government transfers, and regulation of private markets. A large theoretical literature has examined the role and motivation for in-kind subsidies in different sectors of the economy, while a substantial theoretical and empirical literature has studied the supply-side effects of government regulation. Laffont and Tirole (1993) gives a classic reference on the multitude of theoretical issues. Surprisingly, the has been little empirical analysis at the intersection of these literatures, even though there is a growing number of settings where in-kind subsidies affect the decisions of private strategic firms rather than individual consumers. Our paper contributes to the nascent literature at this intersection including: the early work by Gruber and Washington (2005) and Cutler and Reber (1998) on tax and employer subsidizes on employer-sponsored insurance plans; conceptually-related work on government procurement in health care studied in Duggan (2004) and Duggan and Scott Morton (2006); as well as the more recent concurrent work on Medicare Advantage and the Affordable Care Act in Enthoven (2011), Frakt (2011), Decarolis (2015), Curto et al. (2015), Jaffe and Shepard (2018), Tebaldi (2017), and Polyakova and Ryan (2018). We contribute to this literature by analyzing the context of Medicare Part D, which in itself is a large market, as well as by moving the literature forward by exploring ways of incorporating non-constant marginal costs into equilibrium pricing analysis with subsidized prices.

Our paper contributes to the growing literature that analyzes the Part D program as a prominent example of introducing consumer choice in health insurance. On the demand side, a number of papers have explored the rationality of individual choices, consumer myopia, and inertia.[1] A handful of existing studies on the supply side have considered the quantity and quality of the plan menu offered by Part D insurers and price responses to consumer inertia (Lucarelli et al., 2012; Ho et al., 2015; Wu, 2016; Miller and Yeo, 2015; Chorniy et al., 2014; Einav et al., 2016; Fleitas, 2017; Starc and Town, 2015). While studying a different question—the difference in cost-sharing design between independent and integrated drug plans—Starc and Town (2015) is the closest to our work in terms of the modelling approach, including the structural estimation of demand for Part D plans by consumer risk type.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 lays out the empirical model of supply and demand, while Section 4 presents model estimates. Section 5 discusses the economics forces in counterfactual subsidy mechanisms. Section 6 concludes.

[1]Including 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); Bundorf et al. (2013); Ketcham and Simon (2008); Ketcham et al. (2015); Ericson (2014); Miller and Yeo (2014); Abaluck and Gruber (2016); Ho et al. (2015); Polyakova (2016); Heiss et al. (2016); Bundorf et al. (2018); Einav et al. (2015); Abaluck et al. (2015); Gowrisankaran et al. (2015).

2 Economic Environment and Data

2.1 Medicare Part D Primer

Medicare is a public health insurance program covering elderly and the disabled in the United States. Over 50 million individuals benefit from Medicare, which accounts for roughly \$500 billion in annual budgetary outlays. The program is administered by the Centers for Medicare and Medicaid Services (CMS). Most beneficiaries become eligible for the program when they turn 65 and are automatically enrolled into insurance for inpatient (Part A) and outpatient (Part B) services under the so-called "traditional" fee-for-service Medicare. At this point consumers can make two choices. First, consumers can decide to purchase coverage for their pharmaceutical expenditures that is not included in Parts A or B. Such coverage is provided by private Prescription Drug Plans (PDP) under what is known as the Medicare Part D program. Consumers have a choice of more than a dozen PDP plans in each of the program's 34 geographic markets. Alternatively, consumers may decide to opt out of traditional Medicare altogether and switch to a private Medicare Advantage plan for bundled inpatient, outpatient, and pharmaceutical coverage (MA-PD). MA-PD plans provide a privately administered, but publicly financed alternative to government-run Medicare.

Pharmaceutical coverage for Medicare beneficiaries is the empirical context of our analysis. This drug program (launched in 2006) is a large and rapidly growing market that accounts for about a fifth of overall federal spending in Medicare, i.e. about \$100 billion. Beyond its sheer economic size, this market further plays an important role in policy making, as it has become *the* role model for private provision of publicly funded social insurance. Consumers in Part D bear only a fraction of the program's cost (in total circa 15 percent) due to extensive premium subsidies and risk-equalization programs. The efficiency of the mechanism by which the government sets the premium subsidy is at the heart of our research question, so we describe it in some detail.

To determine subsidies for pharmaceutical coverage, CMS starts by collecting "bids" from insurers that should reflect the full price that an insurer would charge for an average risk beneficiary. The regulator then takes a weighted (by lagged enrollment shares) average of these bids across all Part D plans across all markets. Consumer-facing premium for each plan is set by CMS as 25 percent of the national bid average plus the difference between the plan's bid and the national average. In addition to consumer premiums, insurers collect a payment from CMS that varies across consumers depending on their health risk. CMS assigns consumer i a risk score. To a plan that enrolls this consumer, CMS then pays a subsidy equal to the insurer's bid multiplied by i's risk score net of consumer premium.[2]

For consumers with income under 150 percent of the federal poverty line (known as "LIS," or low-income-subsidy consumers), CMS pays the full premium when consumers are enrolled in a qualifying plan. Further, LIS consumers are randomly assigned to qualifying plans, unless they actively enroll in a plan of their choice. Plans in market m qualify for LIS random assignment and full subsidies if their consumer premiums fall below the average consumer premium in market m. As discussed in detail in Decarolis (2015), LIS random assignment generates a discontinuity in market shares that we account for in Section 3.2.

Table 1 reports key summary statistics for the Part D market. 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, 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 0.5 million enrolled in stand-alone prescription drug plans (PDPs).[3].

Consumers had on average a choice of 49 Part D PDP plans in their markets, offered by 16 insurers in 2007-2010. We can clearly see the central role of subsidies in this market: the national average bid in years 2007-2010 was \$1001, \$648 of that amount was covered in subsidies. Consumers paid the remainder (plus any additional premiums that insurers can collect for coverage enhancements) for an average consumer-facing premium of \$505. Consumer premiums varied substantially across geographic markets and time, ranging from \$375 to \$643 in annual premiums.

In our empirical analysis, we differentiate consumers by their health risk type. The idea is that consumers of different health may have different preferences for pharmaceutical coverage and also generate different costs for insurers. We distinguish six consumer risk types. Among consumers not eligible for low income support, we construct five risk groups that differentiate across low risk (relatively healthy) and high risk (relatively unhealthy) consumers. We treat consumers that are eligible for low income support as a separate (sixth) risk group. In Panel D, we report the share of each consumer risk type among enrollees in our analytic sample. The shares of risk type 1 (lowest risk) to 5 (highest risk) consumers in the market are on average respectively: 5, 20, 39, 7, and 1 percent. The remainder 30 percent of potential consumers are individuals eligible for low income subsidies.

[2] In addition to premium subsidies, CMS further provides additional payments for especially high risk consumers. We discuss more details in Appendix A.

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

2.2 Data

We combine two primary sources of data. The first dataset contains detailed information about plan prices and characteristics for all Part D plans in all markets in years 2007 to 2010. The data also includes information on market-level aggregate enrollment in PDP, MA-PD as well as other types of Part D programs.[4] The second dataset are administrative individual-level pharmaceutical claim records for years 2007-2010 for a 20 percent sample of Medicare beneficiaries.[5] These data contain individual-level information on consumer demographics, including chronic conditions, as well as Part D enrollment status, including plan choice and low-income subsidy eligibility. The enrollment data is linked to claim information that records each drug purchase for each consumer in the sample. The purchase records include information about the total cost of prescription as well as how this cost is split between consumer, insurer, and the government. From the 20 percent sample we construct our analytic sample by restricting the data to individuals living in 50 US states and not having special types of pharmaceutical insurance, such as, for example, employer-provided coverage. The restrictions (described in more detail in Appendix Section A) decrease the sample size from 38,628,624 individual-years (for 11,266,409 unique individuals) in the raw data to 23,957,330 individual-years (for 7,543,722 unique individuals) in our analytic sample.

3 Model

We propose an empirical model of demand and supply of insurance contracts in Medicare Part D that will help us evaluate the efficiency and allocative properties of the subsidization mechanism in this 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) before turning to a model of supply that incorporates the many institutional features of this market.

3.1 Demand

We model consumers in 34 Part D markets in years 2007 to 2010 as choosing an insurance plan that maximizes their indirect utility as a function of both pecuniary and non-pecuniary

[4]These data are 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.

[5]Detailed description of data and data access is available in the online supplementary materials for this paper as well as at www.resdac.org.

plan characteristics. We estimate demand separately for six different risk types of consumers. The underlying utility structure is assumed to be the same across all six consumer groups; for the LIS market we adjust plan characteristics to reflect the differences in premiums and cost-sharing that these consumers face.[6]

The utility for enrollee i of plan j in market m consists of a deterministic component and an idiosyncratic Type I Extreme Value-distributed random shock, ϵ_{ijm}:

$$u_{ijm} = -\alpha_i p_{jm} + \beta_i x_{jm} + \xi_{jm} + \epsilon_{ijm}, \tag{1}$$

where p_{jm} is the plan's enrollee-facing premium after subsidies. The observable characteristics, x_{jm}, include 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 quality characteristics of plans, such as pharmacy networks. ξ_{jm} is a plan-specific fixed effect that captures unobserved plan quality. We also include the number of years the plan has been on the market as a reduced-form approach to capturing stickiness in consumer decision-making.[7] The utility of the outside option is normalized to zero. For all five risk types of regular consumers, the outside option constitutes buying a Medicare Advantage drug plan bundled with a medical plan or not buying any drug coverage. For LIS consumers that are randomly assigned to plans when first entering the program, we assume that the outside option constitutes switching to Medicare Advantage.[8]

We model heterogeneity in preferences along two major dimensions. First, for five risk types of regular enrollees, unobserved consumer heterogeneity enters the model through random coefficients on the premium, coverage in the gap, and overall inside option. The

[6] One empirical challenge specific to the LIS market 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 premiums, zero deductibles, no gap in coverage, and otherwise reduced or eliminated cost-sharing.

[7] Additional details and a discussion of the vintage measure may be found in Section C of the Online Appendix.

[8] We do not include consumers eligible for other sources of Part D coverage into our model, assuming that their coverage options are always superior to the publicly available Part D contracts. This primarily includes consumers eligible for coverage through Veteran Affairs and employer sponsored plans.

unobserved heterogeneity may capture differences in income, as well as individuals' differences in risk conditional on risk type, and risk aversion. We choose a log-normal distribution for random coefficients on premiums; it 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:

$$\ln \alpha_i = \alpha + \sigma_\alpha \nu_i. \tag{2}$$

The parameters governing coverage in the gap, β_{gap} and σ_{gap}, and the inside option, β_{inner} and σ_{inner}, are specified analogously without the logarithmic transformation.

Second, we allow the entire vector of preference parameters to vary based on observable risk type (and LIS status) of the consumer. This approach allows for risk-based sorting of consumers in the market. We place non-LIS consumers into one of five groups on the basis of a one-dimensional risk score, which we construct from a normalized prediction of individual-level drug expenditures. We predict drug expenditures using a linear link between historical drug expenditures observed in the data and individual's health status measured with indicators for the presence of more than 50 chronic conditions. We divide individuals into five risk groups using percentiles of the risk score distribution (5th, 25th, 50th, 75th, and 95th percentiles).[9]

3.2 Supply

Modeling the supply side in Medicare Part D market presents a considerable challenge, as the decision-making of insurers is affected by a complex set of regulatory provisions. For brevity, we have relegated the intricate details of the construction of our profit function to Section D in the Online Appendix. The resulting profit function for a given firm f with a portfolio of PDP plans J_f in a given market is given by:

$$\pi_f(b) = \sum_{j \in J_f} \left(\sum_{t=1}^{5} \left[M_t^R s_{jt}^R(b)(\theta_t^R b_j - \kappa_t^R c_j) \right] + M^{LIS} s_j^{LIS}(b)(\theta^{LIS} b_j - \kappa^{LIS} c_j) \right), \tag{3}$$

where b is the vector of bids set by firms, M is market size, s is a vector of plan shares, and c is a vector of marginal costs; the subscripts j refers to each plan and t to each consumer risk type among 5 risk types of regular enrollees. The superscripts refer to regular enrollees (R) and LIS enrollees (LIS).

[9]Section B in the Online Appendix outlines the details of the OLS prediction, risk score construction, and division into five risk groups.

Several key aspects of this profit function differ from a profit function in a standard product market. First, while firms are required to set one bid per plan, b_j, both per-enrollee revenues and costs are indexed by risk type and LIS eligibility. For a consumer of risk type t, a firm that bid b_j receives $\theta_t b_j$, where θ_t is a risk-adjustment factor based on the enrollees' risk score. The risk adjustment system is calibrated so that an average enrollee has a risk score of 1, which implies that firms receive more than b_j for higher risk enrollees and less than b_j for lower risk enrollees. Further, the firms faces a different expected cost for each potential consumer, depending on the consumer's health conditions, pharmaceutical needs, and insurance plan design. To model this aspect of the insurance market, we assume that marginal costs vary multiplicatively by consumer risk type, so that the insurer's marginal costs are given by $\kappa_t^R c_j$, where the scaling factors κ_t^R and κ^{LIS} measure differences in average costs across consumer types.[10]

This profit function thus captures the key difference of insurance markets compared to regular product markets: an insurer's marginal cost curve is a function of all prices in the market due to consumers of different risk types sorting across plans. The slope of the marginal cost curve can be negative (adverse selection) or positive (advantageous selection); our model does not impose a restriction on the direction of selection. Since insurers cannot directly price discriminate in this market, the pooling of risks leads to cross-subsidization across consumer types, with lower risk types subsidizing higher risk types.

Second, consumers face subsidized premiums; the premium formula for regular enrollees is:

$$p_j^R = \max\left\{0, b_j - \bar{b} + \zeta\bar{b}\right\}, \tag{4}$$

where $\bar{b}$ is the enrollment-weighted average bid across all Part D plans in the entire US—critically, this includes not just PDP plans but also MA-PD plans—and ζ is the share of the average bid *not* covered by the baseline federal subsidy.[11] The adjustment ζ is set every year by CMS and is governed by fiscal considerations and the Part D statutes. For example, in 2010, this number was 0.36. Notably, this premium subsidy structure distorts both the absolute and relative prices of Part D plans.

Third, the share of LIS enrollees in plan j is complicated by that market segment's random assignment mechanism: only plans with a consumer premium below the average premium in the region qualify for random assignment of LIS consumers; all other plans receive zero LIS enrollees, unless these enrollees opt out of the random assignment and actively choose these

[10]The construction of these scaling factors is discussed in more detail in Section D in the Online Appendix.
[11]Given that $\bar{b}$ is determined by over 1,500 plans, we assume that firms treat it as a fixed constant.

plans. Section D in the Online Appendix outlines how we address the resulting non-linearity in the plan's market share in the LIS market.

We use the profit equation above and the behavioral assumption that insurers in this market engage in Bertrand price competition to infer plan-type-specific marginal costs and to solve for market equilibria under counterfactual subsidy mechanisms.

3.3 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 government subsidies (G) including the deadweight loss associated with taxation needed to fund the subsidy payment:

$$W = CS + \Pi - \lambda G, \tag{5}$$

where λ is the social cost of raising public revenues. All objects in Equation 5 account for opportunity cost: consumer surplus is measured against the zero-utility outside option, profits are computed against what the firm could have made selling to consumers in another market, for example, MA-PD; and government expenditures reflect the opportunity cost of subsidizing the consumer in another market, such as MA-PD or other public pharmaceutical programs. When solving for a vector of prices that would lead to a socially optimal allocation, we adjust the welfare function by multiplying Π by λ. This captures the idea that under a social planner's allocation, the government directly controls prices and will tax/subsidize firms to achieve a zero profit condition. Detailed derivations of each component of the welfare function can be found in Section E of the Online Appendix.

4 Model Estimates

4.1 Demand Parameters

Table 2 reports demand estimates. Columns (1)-(5) report estimates for regular enrollees from the random coefficient logit model with a log-normally distributed price coefficient; column (6) reports demand estimates from a Berry (1994) logit model for LIS enrollees.[12] All models are estimated using instrumental variables to account for the possibility that there is

[12] We do not estimate random coefficients for the LIS enrollees as the relevant product characteristics are set to zero for this population.

unobserved quality aspect of plans in the error term that is correlated with premiums, but that we fail to capture with the observed characteristics. 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). We use BLP-style instruments that measure the number of insurance contracts offered by the same insurer in a different market, as well as Hausman-style instruments that measure prices charged for the similar plans in other geographic markets.[13] These instruments are 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 in prices that is common across these contract (e.g. due to insurer's price negotiations with pharmaceutical producers), but is not correlated with market-specific unobserved quality (e.g. due to local marketing) over and above average quality captured by insurer fixed effects. The first stage is jointly statistically significant with an F-statistic of 245 for the market with regular consumers and 23 for the LIS market.

We find intuitive patterns for the price coefficients, with riskier types and LIS consumers having generally lower price sensitivity than other consumers. We do not find evidence of statistically-significant dispersion in the price coefficient, which is reasonable given that we are estimating demand within groups of consumers with similar expected costs and health risks, both likely drivers of price sensitivity. Modal aggregate elasticities of demand by consumer risk group are -12.9, -8.9, -5.5, -5.26, and -5.9, in the order of increasing health risk. These are economically reasonable estimates and are similar to the range of (aggregated across risk types) elasticities reported in Lucarelli et al. (2012) (-2.0 to -6.0) and Starc and Town (2015) (-5.0 to -6.3).

Non-premium plan characteristics are estimated to have coefficients with intuitive signs. Consumers dislike higher plan deductibles—more so if they have lower health risk—but enjoy measures of plan generosity: coverage in the gap, broader coverage of common drugs, and more in-network pharmacies 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 find some evidence of significant dispersion in preferences (among some risk types) for two other

[13] 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.

variables for which we allow random coefficients: the inner option and the dummy for gap coverage.

To assess whether the estimated willingness to pay is reasonable consider the following calculation. A typical plan with a non-zero deductible has a deductible that varies from \$265 to \$310 dollars. At the average of \$290, removing the deductible has a dollar value of $-290 * (-10.6/23.5) = \$130$ at the median value of the premium coefficient for risk group 1. For risk group 5, the same computation suggests that this group values removing the deductible at $-290 * (-7.92/12.9) = \$178$. These levels of willingness to pay for zero deductible seem reasonable. Noting that consumers still on average pay Medicare's standard 25 percent co-insurance for the first \$290 in spending, the expected monetized value of going from \$290 to zero deductible is $0.8 * (\$290 - 0.25 * \$290) = \$174$, which is almost exactly our willingness to pay estimate for consumers who are likely to spend through the deductible (i.e. those in risk group 5) and it is lower for consumers who are less like to spend the whole amount (the deductible level lies roughly at the 20th percentile of the spending distribution - see Einav et al., 2016).[14]

Column (6) reports 2SLS estimates of the Berry logit model for the LIS market. To estimate LIS demand, we adjust premiums to reflect LIS-specific subsidies and remove cost-sharing rules such as a deductible that LIS consumers do not face. The estimated price coefficient at -7.9 suggests that LIS demand is less sensitive to prices than all risk types of regular enrollees. This is intuitive, as prices are about \$400 lower per year for the LIS enrollees when compared to regular premiums.

4.2 Marginal Cost Estimates

We address two challenges in constructing the marginal cost estimates. First, plans that are eligible and compete for random assignment of low-income subsidy beneficiaries have a non-linear share function preventing us from using a standard approach of inverting the

[14] A similar calculation for coverage in the gap gives a valuation of $\$1000 * 4.62/12.9 = \$358$ for risk group 5 and $\$1000 * 3.83/23.5 = \$163$. Again, this is reasonable given that consumers are not very likely to enter the gap. Further, even if offered, gap coverage will not cover 100 percent of expenditures, making the actual difference between having and not having coverage for the \$3000 gap roughly 75 percent of the gap's amount. 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 \$1,125. 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 that is substantially above this amount for the riskiest consumers and substantially below for the lowest risk consumers.

first order condition for these "distorted" plans. Second, marginal costs are assumed to not be constant within a plan across enrollees—we allow for marginal costs to vary across five regular consumer risk types and LIS consumers.

We start with the estimation of marginal costs for plans that we identify as systematically not competing for random assignment of LIS consumers.[15] This set of plans includes 756 out of 1,540 plans available in 2010, which is the year that our counterfactual analysis will focus on. For these plans, we assume that the pricing incentives are captured by the first order condition of the profit function in Equation 3 with respect to bid b. The first order condition for prices can be inverted to recover the baseline marginal cost c_j for each plan j (Nevo, 2001). Normalizing the marginal cost multipliers κ_t^R (which we estimate from claims data as described in Appendix D) so that the multiplier for risk group 1 of regular enrollees, κ_1^R, is equal to one, the inversion recovers the marginal cost for the least expensive risk group 1 enrollees. We then apply κ_t^R multipliers for risk groups 2-5 as well as the LIS consumers to compute marginal costs for all enrollees.[16]

We next proceed to estimating the marginal costs for plans that we hypothesize distort their bids to compete for LIS random assignment, since we observe these plans' insurer-market pair qualifying for random assignment at least once in our data. We use a hedonic regression of marginal costs on plan characteristics of "non-distorted" plans and then apply this projection to the distroted plans to get a prediction of marginal costs. The hedonic regression for 756 non-manipulating plans takes the following form (we estimate a separate regression for each consumer risk type t):

$$mc_{jt} = X_{jt}\beta_t + \tau_{ft} + \delta_{mt} + \epsilon_{jt}, \tag{6}$$

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 (τ_{ft}) and market (δ_{mt}) fixed effects to account for inherent differences in marginal costs across insurers and geographic

[15] 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. See Decarolis (2015) for a discussion of factors that are likely to drive plans to compete for low income beneficiaries.

[16] As the multipliers do not vary by plan, we assume that on average consumer-type spending does not change across plans, i.e. assume no moral hazard. While the literature has documented the presence of some moral hazard in this market (Einav et al., 2015), the (relatively small) estimated magnitudes are concentrated in the coverage in the gap benefit phase, which only few enrollees reach - hence, we would not expect them to affect the average across a large number of enrollees within enrollee types.

regions. Panel B of Table 2 reports the coefficients for the hedonic regression. Intuitively, the most important determinants of marginal costs are estimated to be deductibles (a higher deductible is associated with a lower marginal cost) and coverage in the gap (plans that offer coverage in the gap have higher marginal costs). We use the estimates of how plan characteristics translate into marginal costs to predict marginal costs for all plans that we assume are "distorted" by LIS random assignment. 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 theirs bids, 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 conditional on the same risk pool.

The resulting vector of marginal costs is centered at $663 for the lowest risk enrollees and ranges from $337 to $1,193 across plans. The scaling for highest risk regular enrollees of type 5 implies an average marginal cost for this risk group of $1,800 with a range across plans from $809 to $3,220. These estimates of economic marginal costs from the inversion of the first order conditions appear plausible given our estimates of plans' accounting costs from the claims data. Our computations suggest that the average PDP plan liability was $588 for regular consumers of risk type 1; $1,067 for average risk category 3; $1,977 for risk category 5, and $1,363 for LIS beneficiaries.[17] Our estimates of economic marginal costs imply profit margins of 7 percent on average (standard deviation of 9 percent) for regular enrollees. These are fairly low margins, suggesting that the regular enrollee market is reasonably competitive, which is consistent with the policy analysis of this market (Congressional Budget Office, 2014).

4.3 Measuring Government Spending

In the remainder of the paper, we repeatedly calculate welfare that requires several assumptions about the computation of government spending. There are two types of government expenditures that we compute. First, we compute premium subsidies. The baseline premium subsidy is set at circa 70 percent (with some minor annual variation) of the average (weighted) bid for basic coverage across all plans offered in the US in a given year. The

[17] We define plan liability from the claims data as follows: for each individual we take the difference between the total cost paid for drugs at the point of sale, subtract patient cost-sharing payments, LIS cost-sharing subsidies, and 85 percent of spending in the catastrophic part of the benefit, as the plan carries only 15 percent liability in the latter benefit phase.

consumer premium is computed as the difference between plan's bid together with plan's add-on prices for any coverage enhancements and the baseline subsidy. The baseline subsidy is not actually paid out to the plans. Instead plans receive a payment that is the difference between their baseline bid multiplied by the enrollee's risk score, and consumer premium. As the risk score can be smaller than 1 for relatively healthy enrollees, plans receive subsidy payments that are much lower than 70 percent of the average (weighted) bid for less risky enrollees, while receiving much higher subsidies for higher risk enrollees. As detailed above and in the Appendix Section B, we construct proxies for risk scores from the information about chronic conditions and use the average risk score per consumer type to scale subsidies received by plans.

The second type of government payment to Part D plans is reinsurance. This payment covers 80 percent of prescription costs for very high spending beneficiaries. In 2010, beneficiaries had to spend more than $6,440 in total on drugs for the re-insurance program to start paying out to plans. We compute average differences in reinsurance payments per risk type (most of the payments are concentrated in risk group 5) from observed claims data. We assume that this payment multiplier is fixed and does not change across counterfactuals. We apply these multipliers to plan-level reinsurance statistics reported by CMS.[18] We further compute average premium and reinsurance subsidies in the MA-PD program, as well as additional payments for low income beneficiaries from the micro-level claims data. The MA-PD computation allows us to estimate the opportunity cost of government expenditures in the PDP program. In our data, most individuals that switch out of PDP plans switch to MA-PD rather than to no coverage. Hence, in our counterfactuals, 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. The details of these calculations are outlined in Appendix Section E.

4.4 Efficiency of the Observed Subsidy Mechanism

Using demand and marginal cost estimates, we next compute consumer surplus, producer profits, government transfers, and total surplus for the observed market allocation and the observed subsidy mechanism. For expositional clarity, we report results for a single year (2010). The calculations are reported in the first and second columns of Table 3 for regular and LIS enrollees, respectively.

[18] Reported in www.cms.gov/Medicare/Medicare-Advantage/Plan-Payment/Plan-Payment-Data.html.

We estimate that total annual consumer surplus generated by Part D PDP for regular enrollees was \$2.4 billion (row 1), or about \$300 for each of 7.8 million enrollees (36 percent of potential market size). The majority of enrollees, 59 percent, were of about average risk, in risk group 3. 6 percent were in the healthiest category and only about 1 percent of enrollees fell into risk type 5. All enrollees paid on average \$510 in premiums, while firms collected \$1,129 in per capita revenue prior to risk adjustment. Accounting for risk adjustments, but not counting any ex post risk corridor payments, we compute that insurer profits amount to \$536 million (row 2). The \$3 billion of consumer and producer surplus came at a steep price—government expenditures on PDP subsidies (including premium subsidy, risk adjustment payments, and reinsurance) totaled nearly \$5.5 billion (row 7). Taken at face value, the program thus generated negative surplus with a return of -46 cents on a dollar spent in subsidies (row 13).

To interpret this computation, however, it is important to take into account the opportunity cost of government funds. The outside option in our model includes either purchasing an MA-PD plan or not purchasing any creditable Part D coverage. If all consumers that were to leave PDP plans enrolled in MA-PD plans, the government would incur a very similar level of expenditures on these consumers. In rows (8)-(10) we compute that the government would have spent \$5.9 billion if PDP consumers enrolled in MA-PD.[19] The difference between row (7) and (10) that amounts to \$450 million, is the extra government spending on pharmaceutical coverage generated by the PDP program. This extra spending 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 on PDP plans of \$585 million.[20]

Putting it all together, we estimate the total surplus generated by the regular Part D PDP market when accounting for the opportunity cost of government funds was about \$3.5 billion. In other words, the government generates extra 65 cents of surplus for each dollar it spends in the PDP program. This positive return on a dollar is one of our primary findings, along with the corollary that the vast majority of this surplus comes from foregone government expenditures. The latter generalizes to many publicly funding settings, where

[19] This estimate of alternative government spending only includes the subsidies for pharmaceutical coverage in MA-PD. The literature on Medicare Advantage has further estimated that there are differences in per capita public spending on *medical* insurance between enrollees in traditional fee-for-service Medicare versus Medicare Advantage. We do not take into account this difference, as we do not have the data to compute risk-type specific differences between traditional Medicare and Medicare Advantage. Since the literature estimates that Medicare Advantage leads to higher government spending (Curto et al., 2015, estimate that the difference was about 3% in 2010 and 11% on average in years 2006 to 2010), our computation is a lower bound of surplus in the PDP program.

[20] We examine the sensitivity of our results to the value of λ in Appendix E.

the return on public spending in each program is hard to evaluate in vacuum, as there are almost always substitute programs where the government still incurs expenditures on the same individuals. Without taking into account "competing" programs and the opportunity cost of government funds, one may significantly underestimate the surplus generated by each publicly-funded program.

Column (2) of Table 3 reports a similar calculation for the LIS market. Consumers in this market enjoyed \$2.6 billion in surplus. We do not report the profits associated with this part of the market, as the the static Bertrand-Nash model of competition used to recover marginal costs in the regular market does not apply to firms engaged in dynamic competition for LIS enrollees. Computing government subsidies and government opportunity cost for LIS enrollees requires some additional accounting 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 we compute from the claims data. These generate significant quantitative changes to subsidy levels as compared to regular beneficiaries—we compute the per capita government spending on LIS enrollees in 2010 to be \$5,287 as compared to \$698 for regular enrollees. 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 negative \$6.1 billion. This is driven by two factors. First, LIS beneficiaries face much lower prices—on average \$485 lower annual premiums than regular enrollees - this difference is paid from public funds. Second, given our computation on the level of subsidization of LIS enrollees in both PDP and MA-PD programs, MA-PD appears modestly less expensive. Hence, the net government spending component of the welfare function for LIS enrollees is negative. The bottom line is that the government spends an enormous amount of money on LIS enrollees in PDP, where willingness-to-pay is low and the opportunity cost of government spending is negative.

5 Counterfactual Subsidy Mechanisms

We are interested in understanding how insurer incentives and consumer demand interact with the subsidy mechanism to determine market outcomes. We initially consider a set of counterfactual subsidization mechanisms where the outside option (MA-PD plans) is held constant while we adjust the subsidy mechanism in the PDP market. This conceptual exercise allows us to cleanly illustrate and tease apart the complex economic forces at work. We then turn to a set of counterfactuals where we adjust the outside option to reflect equilibrium changes in the PDP market. The idea is that any changes to subsidies in the PDP market are likely to be mirrored in the MA-PD market—the counterfactuals with adjusted outside option allow us to simulate such parallel changes.

In both settings, we consider two distinct types of subsidy mechanisms. The observed mechanism sets subsidies as a function of the bids submitted by insurers. With a sufficiently competitive product market, the appeal of this approach is that subsidies are linked to the costs of providing the good. This has the advantage of protecting consumers from the risk of cost increases as well as giving policy makers a practical starting point for determining subsidy levels. The downside of such approach is that strategic firms can internalize the fact that the subsidy is more generous when bids are higher, leading to higher profits at the expense of taxpayers. To evaluate this approach to subsidy determination, we consider several local alterations of the existing mechanism. We start by investigating the equilibrium effects of cross-market ties—under the observed mechanism the subsidy in the PDP market depends on insurer bids in MA-PD and LIS markets; we remove these ties in our counterfactual simulations. We then investigate the role of market power for such bid-based subsidy determination by simulating the allocations under the two extremes where every plan is a firm and where all plans are owned by one firm. Lastly, we consider what happens if bid-based proportional subsidies are tied directly to firm bids rather than to a weighted average of such bids as in the observed mechanism.

We then proceed to consider an entirely different type of subsidy—a flat voucher that is not linked to any contemporaneous insurer behavior. Conceptually, this type of subsidy can generate high-powered incentives to lower prices when markets are sufficiently competitive. At the same time, it imposes a greater informational requirement on the government to arrive at a subsidy level and places the incidence of program cost risks onto consumers. Multiple proposals for reforms in Medicare and other publicly subsidized programs envision flat subsidies; hence, evaluating the benefits and drawbacks of such high powered mechanisms is particularly relevant for understanding the proposed policy-making in and outside of Medicare.

Our intent across all counterfactuals is to both understand the specific welfare consequences of different mechanisms in PDP, which is of independent interest given its popularity and the large amount of government expenditures flowing through it, but also to draw out more general lessons about why different mechanisms generated more or less surplus that may be useful in guiding subsidy mechanism design in more general contexts.

5.1 Results with Outside Option Held Fixed

Table 3 shows the results when the outside option is held fixed. We start with two counterfactuals that remove cross-market links across PDP, MA-PD, and LIS markets in columns (3) and (4). Removing LIS market incentives from the regular market leads to an increase in bids and the generosity of the subsidy relative to the observed allocation. In this counterfactual, the firms have no incentive to compete to be below the average premium in order to be eligible for randomly-assigned LIS enrollees. The resulting PDP enrollment of regular enrollees is higher than under the observed allocation, as is the government spending on both premium and reinsurance subsidies (by $1,743M). Total surplus that accounts for the opportunity cost of government funds increases ($230M) despite the fact that additional government spending exceeds the increase in consumer surplus ($380M) and producer profit ($503M); the positive difference in total surplus is driven entirely by the slower change in the foregone government spending in the MA-PD program ($1,241M). Rows (18)-(22) highlight the role of risk selection in the market. Higher subsidies lead to a very slight increase in the share of enrollment from lower risk consumers; the direction of this effect is intuitive and the small magnitude of changes are consistent with the previous literature that has found limited risk screening on prices in this market (Polyakova, 2016). Column (4) reports the outcomes when the average bid in the formula for determining the subsidy in the regular market no longer includes MA-PD plans (which tend to have lower bids than PDP plans). In this case, the baseline subsidy is set as 68 percent of the average Part D PDP bid. Given higher PDP bids, this change leads to another increase in the subsidy relative to column (3). As a result, consumer surplus and profits increase again; this increase, however, is offset by the growth in opportunity-cost-adjusted government spending, which leads to a slightly lower total surplus. The results in column (4) give us a benchmark (simulated within the model) for the analysis of alternative subsidy mechanisms for the regular market.

These results illustrate several general economic forces at play in this market. First, with a highly-subsidized substitute available (i.e. switching to MA-PD plans that offer both medical and pharmaceutical coverage), consumers' baseline willingness-to-pay for plans in the PDP market is very low. Column (10) reports results when the premium subsidy is set to zero; PDP enrollment drops to near zero as consumers leave the PDP market.[21]

[21] The idea of low willingness to pay for health insurance coverage in the presence of even more distant substitutes—such as charity care—has been documented in other settings. For example, Finkelstein et al.

Second, increasing the generosity of the subsidy bolsters total consumer welfare, but does so by giving costly transfers to inframarginal consumers (and changing relative prices within the PDP market, which can lead to further allocative distortions) while also attracting marginal consumers with decreasing valuations for the product. This is reflected in consumer surplus per enrollee: under the observed mechanism it is \$441, without LIS enrollees this number drops to \$377, and without links to the MA-PD market it declines further to \$322.

Third, the sign of welfare outcomes is largely driven by the opportunity cost of government spending. Government expenditures in PDP per enrollee increase from \$698 under the observed allocation to \$738 and \$760 across the two counterfactuals. Even with producer profits included, these costs substantially exceed the benefits that they generate within the PDP market. It is only when the opportunity cost of spending in the MA-PD market is accounted for, does the overall welfare of the PDP market become positive.

Fourth, the government's difficult balancing act on the consumer side of the market—setting premium subsidies to induce the optimal level of sorting across the inside and outside options and across plans within the inside option—is further complicated by possible strategic behavior by insurers. To help assess the degree of market power and strategic markups in this market, Columns (5) and (6) report results for two extreme counterfactuals on the supply side: one where all plans are independent firms, and one where all plans are owned by a single monopolist. In both cases we let the subsidy rule follow the same mechanism as in column (4). There is a substantial degree of market power possible in this market: the average enrollment-weighted bid under the monopolist is over \$79 higher than with atomistic firms, an increase of just under 7 percent. This leads to higher profits and lower consumer surplus. Interestingly, total welfare is nominally lower under both the atomistic firms counterfactual and under a monopolist than in the reference column (4). Atomistic plans attract slightly too many consumers to the market, given how much it costs the government to service them, while the monopolist has too few.

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, such as a standard government-run program or a regulated monopolist, as we do not know anything about the comparative costs of delivering

(2017) find low willingness to pay for health insurance among low income adults in Massachusetts.

similar insurance plans in such scenarios, and therefore we cannot make any claims about additional efficiencies introduced by competition. Also, while we take the marginal costs of firms as given here, it is possible 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 limit our conclusion to the observation that the current ownership structure gives results similar to that if all plans were independent firms.

Fifth, subsidy design matters. We have already shown that linking the level of the subsidy to other markets can have large effects on market outcomes, and we now show that how insurer bids are translated to consumer premiums can have an even more dramatic effect. We solve equilibrium outcomes under two different approaches: a proportional subsidy and a flat subsidy (i.e. a voucher). Columns (7)-(9) show the outcomes for proportional subsidies where consumers pay 5 percent, 32 percent, or 95 percent of the bid submitted by insurers. This subsidy mechanism produces extreme outcomes, primarily due to the exercise of market power by insurers. In the two cases where consumer are largely shielded from bids, insurers increase bids dramatically. The government is left to cover the large gap between the bid and the consumer-facing premium. In the case where consumers pay 5 percent of the bid, enrollment is nearly universal and government expenditures are nearly a staggering $3,200 per enrollee. Of the increase of $55 billion in government expenditures, firms capture $34 billion, or 62 percent. Only when the consumers pay 95 percent of the bid does the mechanism produce positive welfare numbers, although far below the current mechanism. Notably, the extreme counterfactuals where consumers face almost the full cost of coverage—those in columns (9) and (10)—are the only ones that result in positive nominal surplus without accounting for the opportunity cost of government spending. In these counterfactuals, the government spends far less than it generates in consumer and producer valuation, as only consumers with the highest willingness to pay stay in the program.

Inspecting the formula for determining the consumer-facing premium subsidy in Equation 4 reveals that the existing mechanism is similar to a lump-sum voucher from a consumer's perspective; there is essentially a fixed payment that is applied to each plan in the marketplace. To assess this intuition empirically, we solve for outcomes with vouchers running from $0 to $1500 in $100 increments. Columns (10)-(12) show the outcomes of vouchers with the extremes of $0 and $1500 along with the voucher that generated the highest amount of surplus ($800). Figure 1 illustrates our estimates of total welfare across the range of vouchers considered.[22] Total surplus under the fixed outside option (marked with solid black line) is

[22]We note that our model allows for "soft" exit of plans under very low subsidies. We constrain insurers

positive until \$1,300, peaking at \$800 at a level of welfare that is slightly higher, although comparable, to the existing mechanism. In general, vouchers perform much better than proportional subsidies, largely because they preserve the elasticity of demand on the margin while still allowing the policymaker to influence sorting between PDP and MA-PD. In turn, this leads firms to keep their bids reasonably competitive, minimizing the amount of costly transfers from taxpayers to firms. This is exemplified by the average weighted bid actually being lower at the \$800 voucher than at the extremes. We note that the most generous proportional and voucher counterfactuals have very similar enrollments, but are achieved at vastly different levels of social cost—proportional subsidies lead to average bids of nearly \$4,000 and create an order of magnitude higher social loss.

To put our normative findings into perspective, we perform three benchmark computations of the social optimum in this market. We start by computing the social planner's allocation. We assume that the social planner knows consumer demands and marginal costs (both by type) and can directly set prices, but cannot force consumers to purchase certain insurance plans but must incentivize their choices through plan prices. The detailed results are reported in column (1) of Table 4, while Figure 1 shows the level of welfare achieved by the social planner in comparison to market mechanisms. We find the social planner's solution by solving for a set of plan-specific prices in Equation 5. The social planner increases welfare to \$5.4 billion and sets prices that result in large losses for insurers. To illustrate, Figure 2 plots the resulting changes in premiums compared to the observed mechanism, along with changes in market shares by the highest and lowest risk types, in the California market. Plans are ordered from left to right by increasing marginal cost. There are two broad takeaways: first, the social planner adjusts prices to obtain a general shift in market shares to favor less-costly plans. The majority of plans losing market share are the most expensive, while the single biggest increase in share occurs at a low-priced plan. The second takeaway is that there are substantial differences in where consumers of different risk types move, with the largest difference being that the social planner puts more of the most expensive consumers into cheaper plans. The social planner's solution illustrates that consumers are systematically choosing plans that are too socially expensive and are doing so differentially by risk type which leads to inefficient sorting.

We also consider an alternative scenario that captures some of the intuition of the social

to make weakly positive profits: when subsidies are very low, as in the case of a \$0 voucher, insurers may find it profitable to set very high premiums for some of their plans to induce zero enrollment. This is akin to plans exiting the market in an environment where fixed costs accrue primarily at the insurer rather than at the individual plan level.

Figure 1: Welfare under Counterfactual Subsidy Policies

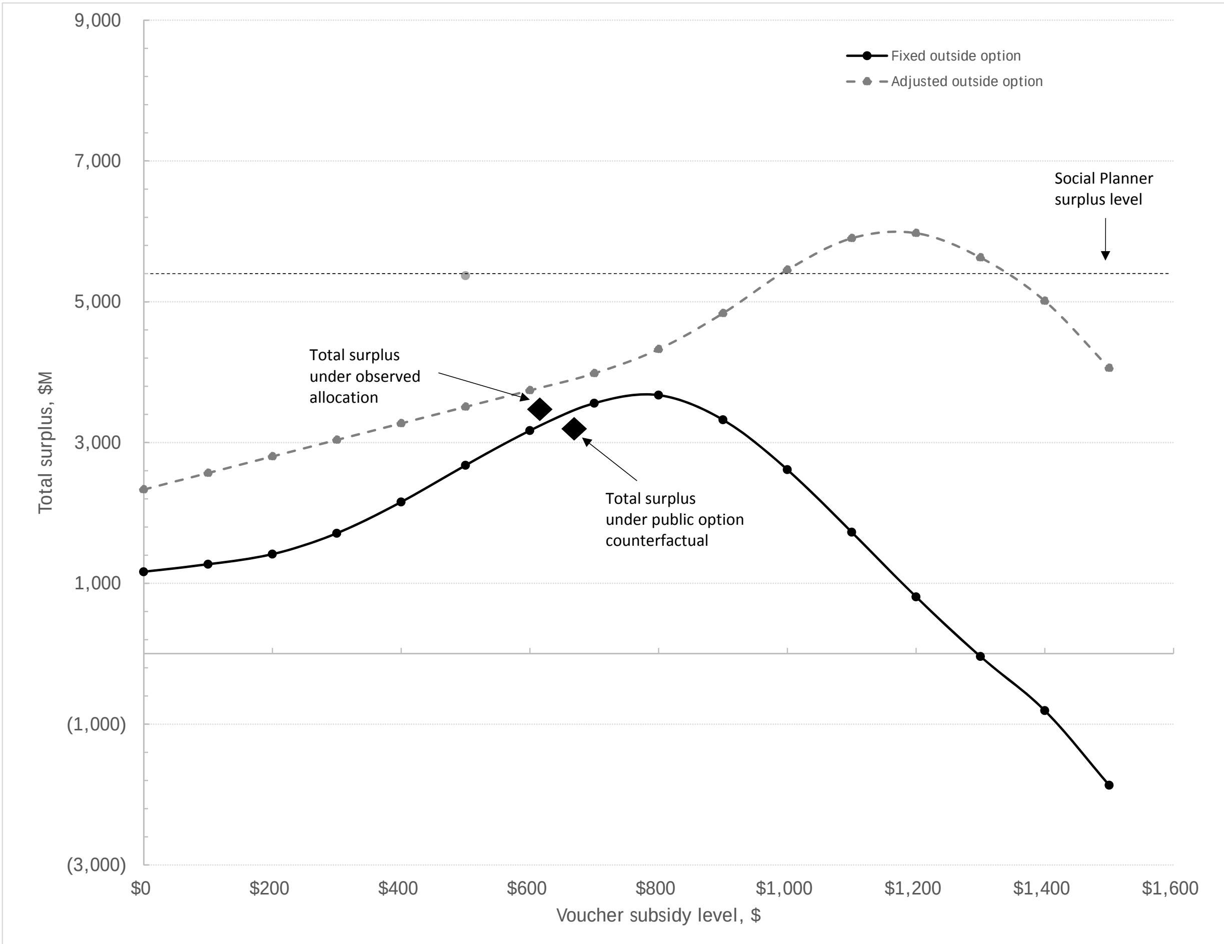

Figure reports estimated total welfare (including the accounting for the the opportunity cost of government spending) in counterfactuals with flat voucher subsidies ranging from $0 to $1,500, in $100 increments. Solid black line marks the welfare estimates in counterfactuals with fixed outside option (row 11 in columns 10-12 in Table 3. Thick grey dashed line plots welfare levels for counterfactuals with vouchers when the subsidy in the outside option (MA-PD) is adjusted to be the same voucher (row 12 in Columns 5-7 of Table 5). We also mark total surplus (at the average subsidy level) for the observed allocation, the public option with subsidy counterfactual, as well as the social planner's surplus level (column 1 row 11 of Table 3; column 2 row 12 of Table 4; and column 1 row 2 of Table 4, respectively).

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, letting Medicaid administer pharmaceutical coverage), we replace all plans in each market with the characteristics of the plan with the lowest estimated marginal cost.[23] We assign the reinsurance subsidies that

[23]We keep the number of plans fixed to equalize the role of the idiosyncratic error term in the logit model

the government pays under the existing mechanism to the marginal cost of that plan. The government then pays the same average premium subsidy as under the observed allocation (\$676) while setting prices such that the firm makes zero profit. The results are reported in Column (2) of Table 4. In this simulation of public option coverage, we find that overall surplus is lower than that under the social planner at \$3.2 billion. Consumer surplus, however, is almost as high as under the social planner at \$3 billion, while producer profit is set to zero by design. This counterfactual manages to achieve a close-to-optimal sorting of consumers between MA-PD and PDP; the total surplus it generates is comparable to the levels of surplus under the optimal voucher and the observed mechanism. This counterfactual is 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. The results in column (3) emphasize that the surplus in this environment is still generated by the opportunity cost of government funds, as consumers do not have sufficient willingness to pay even for the least costly plans. In column (3), we simulate a related counterfactual where the government offers a low cost public option that is, however, not subsidized. This would be closer to traditional Medicare (rather than Medicaid) expanding pharmaceutical coverage and charging the cost of coverage to consumers. In this case, we see similar patterns as in competitive counterfactuals with no subsidies—enrollment drops almost to zero and only consumers with high enough valuation of coverage enroll in the program, generating positive nominal surplus.

We find that most of the mechanisms are very similar in the composition of risk types for enrollees in the inside option. To the extent that risk sorting changes in more extreme counterfactuals, it follows intuitive patterns consistent with the presence of adverse selection in this market. Figure 3 illustrates how the share of high risk consumers (type 4 and 5) changes in the inside option as voucher-based subsidies get more generous. As vouchers increase, leading to lower prices, the share of high risk consumers falls as lower risk consumers start entering the market. This gradient is relatively shallow at very high subsidy levels, suggesting that after a certain threshold, prices are not the first-order drivers of selection.

when comparing outcomes across counterfactuals. For some lowest cost plans we do not observe positive enrollment for all risk types, which leads to a missing estimate of the plan-risk-type-specific fixed effect ξ. The proceed with the public option counterfactuals we had to impute the missing ξ estimates. We proceeded by taking ξ estimates for a given plan for other risk groups and scaling is by the ratio of average differences in ξ's across risk types among all plans for which we were able to estimate ξ's.

Figure 2: Social Planner's Solution: Changes in Premiums and Market Shares

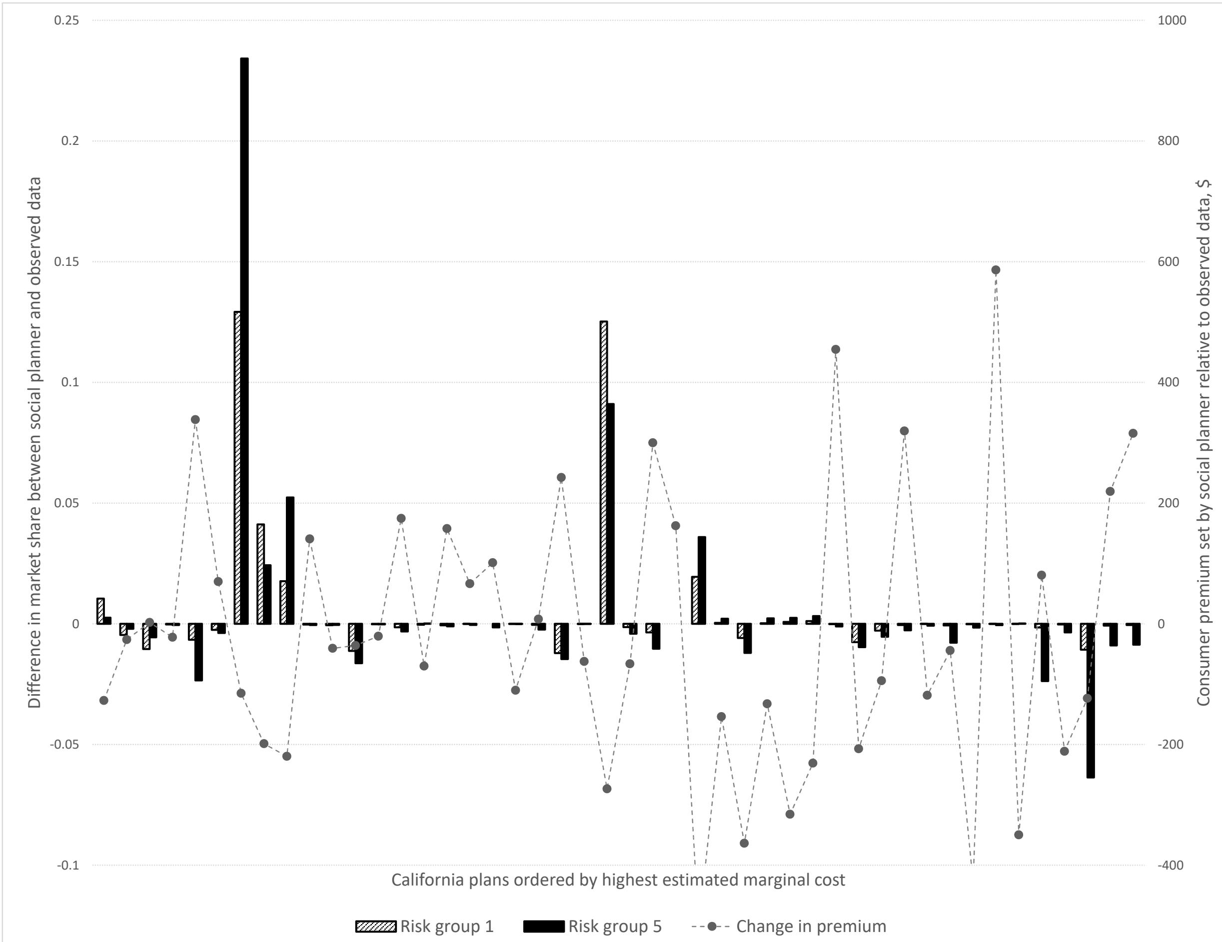

Figure illustrates an example of the social planner's allocation in one Part D market - California in 2010. The x-axis orders 47 PDP plans that were available in California in 2010 by increasing - from left to right - estimated marginal cost. On the left-hand-side y-axis we then plot the change in the market share of each plan that the social planner induces relative to the observed allocation. The changes are plotted with vertical bars. Positive changes imply that the social planner allocates a plan a higher market share, while negative bar values imply a lower market share for the plan, relative to the plan's observed market share. We report the changes in market shares separately for lowest risk regular consumers (risk group 1) and highest risk regular consumers (risk group 5). The right-hand-side y axis and the corresponding dashed line plots the changes in premiums for each plan between the premium set by the social planner relative to the observed premium. A positive value of the premium difference for a plan implies that the social planner's solution would increase the premium for this plan.

For example, moving from \$800 voucher in column (11) to \$1,500 voucher in column (12) leads to a dramatic change in prices from \$444 to \$78—this is associated with the share of inside option enrollment from risk groups 1 and 2 only increasing from 25 to 30 percent, and high risk consumer share decreasing from 15 to 10 percent. Notably, Figure 3 highlights that the optimal voucher of $800 achieves the same share of high risk consumers among inside option enrollement as under the social planner. This optimal share of high risk consumers is about 5 percentage point lower than the observed share, suggesting that too few of low risk consumers are purchasing PDP insurance (a common result in markets with adverse selection).

With the usual caution that our findings are specific to our settings, there are several general themes that can be distilled from our results that may apply to other instances of regulated and subsidized competition. We find that the best mechanisms share three qualities: they preserve the marginal relationship between the prices that the firm sets and the prices that consumers face, they limit how fast the subsidy grows as a function of firm prices, and they link the consumer-facing price to (social) marginal cost. Keeping consumer-facing prices related one-to-one to firm's prices at the margin increases the elasticity of demand and leads to lower prices in equilibrium. Slowing down how fast the subsidy grows relative to the prices that firms set also helps limit firms from exerting market power to increase prices. Finally, as illustrated by the social planner's solution, keeping consumer prices related to marginal cost prevents allocative inefficiency along both the extensive and intensive margins.

The existing Part D mechanism reflects these three qualities in several dimensions. First, linking the equilibrium subsidy with both the LIS and MA-PD markets helps keep the subsidy low. Second, the form of the subsidy, as a percentage of enrollment-weighted average of last year's plans average bid, both limits strategic pricing by insurers and preserves the marginal relationship between bids and consumer prices by acting like a voucher. Third, the size of the subsidy is not so large that all consumer-facing premiums are zero, which in turn helps preserve the link between marginal cost and prices.

5.2 Results with Adjusted Outside Option

We next examine how the results of the main counterfactual simulations change when the outside option (i.e. MA-PD) is adjusted to reflect changes in the generosity of subsidies in Part D PDP. For each counterfactual, we adjust the subsidy in the outside option by the average value of the change in PDP subsidy. Importantly, this remains a partial equilibrium analysis, as we do not consider possible general equilibrium changes in MA-PD prices, in

Figure 3: Risk Sorting under Counterfactual Subsidy Policies

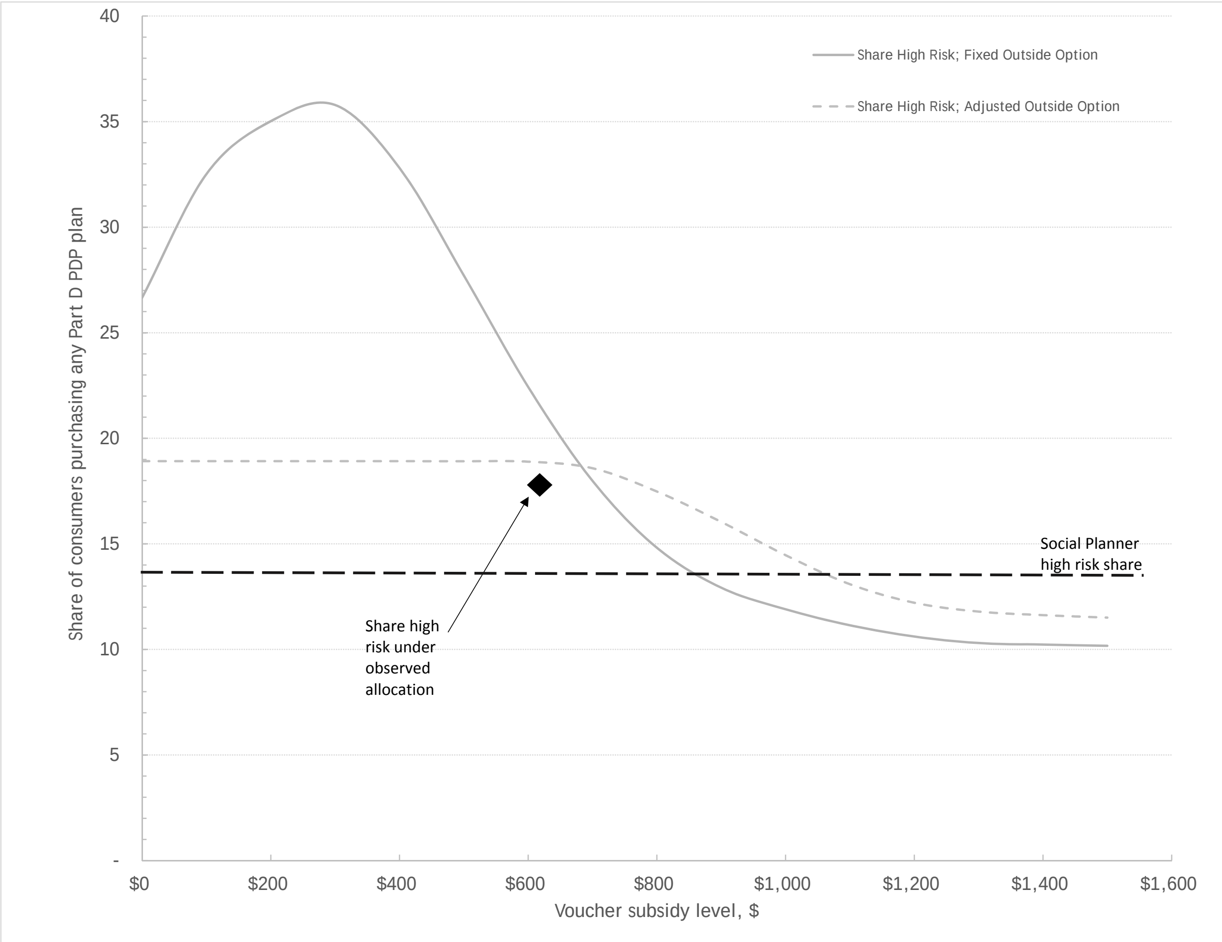

Figure reports the share of high risk consumers (rik type 4 and 5) among regular consumers that are buying the inside option - a PDP plan - in counterfactuals with flat voucher subsidies ranging from \$0 to \$1,500, in \$100 increments. Solid grey line marks the share of high risk consumers in counterfactuals where the outside option is held fixed. Short-dashed grey line marks the share of high risk consumers in counterfactuals that adjust the outside option to have the same level of voucher subsidy as the inside option. Longer-dashed grey line demarcates the share of high risk consumers in the PDP market as simulated under the social planner's allocation. The plotted quantities are also reported in rows 21 plus 22, columns 10-12 of Table 3 (for fixed outside option); rows 22 plus 23 of columns 5-7 in Table 5 (for adjusted outside option) and rows 22 plus 23 of column 1 in Table 4 (for social planner).

prices of other pharmaceutical insurance options, or drug prices.[24] We intend for this adjustment to capture the idea that when a government reforms one publicly-subsidized program,

[24]In practice, this is implemented by adjusting the utility of choosing the inside option. We discuss the details of the adjustment calculation in Appendix F.

it is likely to implement similar reforms in related programs. In the case of Medicare Part D, PDP and MA-PD are obviously closely-related substitutes, but this general idea holds more broadly—few publicly-subsidized program exist in isolation. The results are given in Table 5. There are two additional rows (6 and 26) included in the table in comparison to Table 3, listing the additional payments that the government makes to the MA-PD in each counterfactual relative to the observed subsidy levels in MA-PD (row 6), and the dollar amount of the adjustment we make to the outside option (row 26). A positive number in row (26) implies that the outside option becomes more attractive. There are two ways in which MA-PD related computations change in these results. Row 6 accounts for how much more or less the government needs to pay to consumers that purchase the MA-PD (i.e. outside) option. Rows 8 and 9 that measure the opportunity cost of government funds are also adjusted to reflect a similar change in what the government would have paid for PDP enrollees if they switched to MA-PD.

A common thread across all counterfactuals with adjusted outside option is that sorting between the inside and outside options (and across risk types within the inside option) is very similar to the observed allocation. This is intuitive, as by adjusting the outside option, we are essentially restoring the observed differences between the two enrollment options. In counterfactuals in which we remove the linkages across LIS, MA-PD and PDP markets, or alter the degree of market power, PDP insurer bids and consumer premiums decrease slightly in face of a more attractive and hence more competitive outside option (columns 1-4). Consumer surplus increases while government payments net of the opportunity cost decrease, given that we are subsidizing PDP and MA-PD in similar ways. As a result, total surplus accounting for the opportunity cost of government funds in these counterfactuals is roughly $500M higher relative to that computed for the same countefactuals in Table 3.

We observe similarly intuitive changes in the counterfactuals that consider flat subsidies. As Figure 1 shows, for voucher levels that are close to the observed subsidy level (up until around $1,000) the surplus of counterfactuals with adjusted outside option is reasonably close—only somewhat higher—to that of both the observed allocation and the counterfactuals without the outside option adjustments. The optimal voucher with adjusted outside option is substantially higher, at $1,200 versus $800 in the non-adjusted case. This is again driven by the more attractive outside option creating competition for PDP plans leading to lower bids, premiums, and relative subsidies. Figure 3 highlights the stabilization of the market movement when we adjust the outside option—the share of high risk consumers remains stable at the observed levels until the voucher increases above $900; at the highest voucher levels, the equilibrium level of high risk consumers converges to the same level as that under the counterfactuals without outside option adjustments, remaining somewhat closer to the optimal social planner level. Moreover, as we observe in Figure 1, with high vouchers applied to both the inside and outside options, the market achieves the level of surplus that is equal to the social planner. This effect is generated entirely by significant changes in the opportunity cost of government funds that results when we adjust the outside options that is driven by fixed MA-PD prices that do not respond to more generous subsidies. This highlights the general principle that to the extent that the government can control the level of prices in the market (in this case this happens mechanically as we do not allow MA-PD prices to adjust), higher subsidies can generate substantial consumer surplus. This surplus is still generated at very high nominal government spending. Without accounting for the opportunity cost of government funds, the voucher of $1,200 (column 6) generates negative $13B in surplus, loosing 39 cents on each dollar spent.

The general takeaway is twofold. First, the evaluation of subsidized programs is challenging in the presence of possibly subsidized substitutes, which is commonplace in many settings. Whether there is a positive return on a dollar spent in one program depends crucially on how this dollar would have been spent in related programs on the same beneficiaries. This phenomenon is very transparent in markets with close substitutes such as PDP and MA-PD, but is likely to still be important, but less obvious, in other programs (for example, health insurance coverage and charity care as examined in Finkelstein et al., 2017). Second, if substitute and related programs are likely to be subject to the same policies as the program of interest, estimating the general equilibrium effects may be necessary to understand the full economic impact of changes in both programs. Without general equilibrium estimates, it is useful to focus on partial equilibrium analysis that holds substitutes fixed to understand the economic forces of, for example, the subsidy design mechanisms as we do in Table 3.

6 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 for our specific empirical setting. First, we find that the current PDP program is efficient only if we account for the fact that the government would likely subsidize the same consumers outside of the PDP program 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 two related factors. First, demand for PDP plans is generated almost exclusively by high subsidies—consumers have very low willingness-to-pay for unsubsidized plans, driven by the availability of close substitutes. Second, this market is imperfectly competitive 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 mutes insurers' ability to raise subsidies, and hence positively affects total welfare. This is due to the complex way in which prices for distinct parts of the program, such as Medicare Advantage Prescription Drug coverage, Low Income Subsidies, and market premiums for regular beneficiaries in PDP plans, are all tied together into one mechanism. We find that the current mechanism that incorporates 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 setting proportional subsidies would lead to a rapid upward price spiral, as the competitive pressure on the market is not strong enough to mitigate the "raising-the-subsidy" incentives.

Further, our analysis reveals a close connection between Part D PDP and Medicare Advantage that, although not emphasized in prior 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 the analysis of many other public programs that do not exist in isolation, but, instead, are linked to each other through the choices of consumers and producers or through government transfers.

Beyond the Part D context, our setting that is characterized by the presence of two publicly subsidized programs that are close substitutes, sheds light on the challenges inherent in the analysis of economic returns to any dollar spent on social insurance programs. In many such programs, and especially healthcare, the government faces a version of the Samaritan's dilemma. If public funds are ultimately used to pay for individuals' healthcare needs through *some* channel, then the question the policy-maker faces is not whether to subsidize healthcare use, but finding the most efficient way of doing it. This idea lies at the heart of our results—funding the public benefit that we analyze makes economic sense only insofar as there would be some expenditures on the activities related to this benefit in all counterfactual policies that are plausibly available to the government. This insight has broad implications for the empirical analysis of economic returns to many other public policies that is frequently done in isolation of potentially less obvious substitutes.

While our empirical analysis focused on the subsidy mechanism in the Medicare Part D program, our findings have broader implications for market design of privately-provisioned and publicly-subsidized social insurance programs. As in any setting with the equity-efficiency trade-off, 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 the paper we illustrate that the details of the subsidy mechanism matter dramatically for how these incentives play out. Further, depending on whether the policy is guided by the considerations of consumer surplus, total welfare, or government spending, we demonstrate that different subsidy mechanisms deliver drastically different results across these three measures of surplus. Overall, we argue that, the less-studied supply-side behavior in the presence of regulatory intervention and subsidization plays a key role in determining the efficiency outcomes of privately-provided social insurance programs. Answering the general question about the optimal mechanism design in these increasingly economically relevant settings presents an important avenue for future empirical and theoretical research.

References

Abaluck, J., J. Gruber, and A. Swanson, "Prescription Drug Use under Medicare Part D: A Linear Model of Nonlinear Budget Sets," *NBER Working Paper No. 20976*, 2015.

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101(4)*, 1180–1210.

_ **and** _ , "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance," *American Economic Review*, 2016, *106(8)*, 2145–2184.

Berry, Steve, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63(4)*, 841–889.

Berry, Steven T, "Estimating discrete-choice models of product differentiation," *RAND Journal of Economics*, 1994, *25(2)*, 242–262.

Bundorf, K., M. Polyakova, and M. Tai-Seale, "Delegating Decision-Making to the

Machine: Experimental Evidence from Health Insurance," *Unpublished mimeo, Stanford University*, 2018.

Bundorf, M. K., R. Mata, M. Schoenbaum, and J. Bhattacharya, "Are Prescription Drug Insurance Choices Consistent With Expected Utility Theory?," *Health Psychology*, 2013, *32* (9), 986–994.

Chorniy, A., D.P. Miller, and T. Tang, "Mergers in Medicare Part D: Decomposing Market Power, Cost Efficiencies, and Bargaining Power," *Princeton University Unpublished Mimeo*, 2014.

Congressional Budget Office, "Competition and the Cost of Medicare's Prescription Drug Program," July 2014.

Curto, V., L. Einav, J. Levin, and J Bhattacharya, "Can Health Insurance Competition Work? Evidence from Medicare Advantage," *NBER Working Paper No. 20818*, 2015.

Cutler, David M. and Sarah J. Reber, "Paying For Health Insurance: The Trade-Off Between Competition And Adverse Selection," *Quarterly Journal of Economics*, 1998, *113(2)*, 433–466.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, 2015, *105(4)*, 1547–1580.

Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi, "Do switching costs make markets less competitive?," *Journal of Marketing Research*, 2009, *46(4)*, 435–445.

Duggan, M., "Does Contracting Out Increase the Efficiency of Government Programs? Evidence from Medicaid HMOs.," *Journal of Public Economics*, 2004, *88(12)*, 2549–2572.

_ **and F. Scott Morton**, "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing," *Quarterly Journal of Economics*, 2006, *121(1)*, 1–30.

Einav, L., A. Finkelstein, and P. Schrimpf, "The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D," *Quarterly Journal of Economics*, 2015, *130(2)*, 841–899.

Einav, Liran, Amy Finkelstein, and Maria Polyakova, "Private provision of social insurance: drug-specific price elasticities and cost sharing in Medicare Part D," *NBER Working Paper No. 22277*, 2016.

Enthoven, Alain C., "Reforming Medicare by Reforming Incentives," *New England Journal of Medicine*, 2011, *364(21).*

Ericson, Keith, "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription

Drug Insurance Exchange," *American Economic Journal: Economic Policy*, 2014, *6(1)*, 38–64.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," Working Paper 23668, National Bureau of Economic Research August 2017.

Fleitas, S., "Dynamic Competition and Price Regulation When Consumers Have Inertia: Evidence from Medicare Part D," *Unpublished mimeo, University of Arizona*, 2017.

Frakt, Austin, "The key difference between exchanges and the Ryan-Rivlin plan," www.theincidentaleconomist.com March 2011.

Gowrisankaran, G., C. Marsh, and R. Town, "Myopia and Complex Dynamic Incentives: Evidence from Medicare Part D," *NBER Working Paper No. 21104*, 2015.

Gruber, J., "Delivering Public Health Insurance Through Private Plan Choice in the United States," *The Journal of Economic Perspectives*, 2017, *31* (4), 3–22.

_ **and E. Washington**, "Subsidies to Employee Health Insurance Premiums and the Health Insurance Market," *Journal of Health Economics*, 2005, *24*, 253–276.

Hausman, Jerry A. and James M. Poterba, "Household Behavior and the Tax Reform Act of 1986," *The Journal of Economic Perspectives*, 1987, *1(1)*, 101–119.

Heiss, Florian, Adam Leive, Daniel McFadden, and Joachim Winter, "Plan Selection in Medicare Part D: Evidence from Administrative Data," *Journal of Health Economics*, 2013, *32*, 1325–1344.

_ **and Viktor Winschel**, "Likelihood approximation by numerical integration on sparse grids," *Journal of Econometrics*, 2008, *144(1)*, 62–80.

_ **, Daniel Mcfadden, and Joachim Winter**, "Mind the Gap ! Consumer Perceptions and Choices of Medicare Part D Prescription," in David Wise, ed., *Research Findings in the Economics of Aging*, The University of Chicago Press, 2010, pp. 413–481.

_ **, Daniel McFadden, Joachim Winter, Amelie Wuppermann, and Bo Zhou**, "Inattention and Switching Costs as Sources of Inertia in Medicare Part D," *NBER Working Paper No. 22765*, 2016.

Ho, Kate, Joseph Hogan, and F. Scott Morton, "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program," *NBER Working Paper No. 21028*, 2015.

Hortacsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller, "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market," *NBER Working Paper No. 20988*, June 2015.

Jaffe, S. and M. Shepard, "Price-Linked Subsidies and Imperfect Competition in Health Insurance," *NBER Working Paper No. 23104*, 2018.

Kesternich, Iris, Florian Heiss, Daniel McFadden, and Joachim Winter, "Suit the action to the word, the word to the action: Hypothetical choices and real decisions in Medicare Part D," *Journal of Health Economics*, 2013, *32(6)*, 1313–1324.

Ketcham, J. and K. Simon, "Medicare Part D's Effects on Elderly Drug Costs and Utilization," *American Journal of Managed Care*, 2008, *14(11)*, 14–22.

_ , C. Lucarelli, and C. Powers, "Paying Attention or Paying Too Much in Medicare Part D," *American Economic Review*, 2015, *105(1)*, 204–233.

Ketcham, Jonathan D, Claudio Lucarelli, Eugenio J Miravete, and M Christopher Roebuck, "Sinking, Swimming, or Learning to Swim in Medicare Part D," *American Economic Review*, 2012, *102(6)*, 2639–2673.

Klemperer, Paul, "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade," *Review of Economic Studies*, 1995, *62(4)*, 515–539.

Kling, J. R., S. Mullainathan, E. Shafir, L. C. Vermeulen, and M. V. Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *Quarterly Journal of Economics*, 2012, *127(1)*, 199–235.

Laffont, J.-J. and J. Tirole, *A Theory of Incentives in Procurement and Regulation*, The MIT Press, 1993.

Lucarelli, Claudio, Jeffrey Prince, and Kosali Simon, "The Welfare Impact of Reducing Choice in Medicare Part D: A Comparison of Two Regulation Strategies," *International Economic Review*, 2012, *53(4)*, 1155–1177.

McGuire, T. G., J. P. Newhouse, and A. D. Sinaiko, "An economic history of Medicare Part C," *The Milbank Quarterly*, 2011, *89* (2), 289–332.

Miller, Daniel P and Jungwon Yeo, "A Dynamic Discrete Choice Model of Switching Costs in Medicare Part D," *Clemson University and Singapore Management University Unpublished Mimeo*, 2014.

Miller, D.P. and J. Yeo, "The Consequences of a Public Health Insurance Option: Evidence From Medicare Part D," *Clemson University and Singapore Management University Unpublished Mimeo*, 2015.

Miller, Keaton, "Do Private Medicare Firms Have Lower Costs?," *University of Oregon Unpublished Mimeo*, 2014.

Nevo, A., "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 2001, *69*, 307–342.

Oliver, Thomas R., Philip R. Lee, and Helene L. Lipton, "A Political History of Medicare and Prescription Drug Coverage," *Milbank Quarterly*, 2004, *82(2)*, 283–354.

Polyakova, M., "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D," *American Economic Journal: Applied*, 2016, *8(3)*, 165–95.

_ **and S. Ryan**, "In-kind Transfers, Tagging, and Market Power: Evidence from the ACA," *Unpublished mimeo, Stanford University and Washington University in St Louis*, 2018.

Small, Kenneth A. and Harvey S. Rosen, "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 1981, *49(1)*, 105–130.

Starc, Amanda and Robert J. Town, "Internalizing Behavioral Externalities: Benefit Integration in Health Insurance," *NBER Working Paper No. 21783*, 2015.

Tebaldi, P., "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *Unpublished mimeo, University of Chicago*, 2017.

Vetter, S., F. Heiss, D. McFadden, and J. Winter, "Risk attitudes and Medicare Part D enrollment decisions.," *Economics Letters*, 2013, *119(2)*, 128–132.

Williams, H.C.W.L, "On the formation of travel demand models and economic evaluation measure of user benefit," *Environment and Planning*, 1977, *9(3)*, 285–344.

Winter, J., R. Balza, F. Caro, F. Heiss, B. Jun, R. Matzkin, and D. McFadden, "Medicare prescription drug coverage: Consumer information and preferences," *PNAS*, 2006, *103(20)*, 7929–7934.

Wu, Yufei, "Supply Response to Consumer Inertia: Strategic Pricing in Medicare Part D," *MIT Unpublished Mimeo*, 2016.

Table 1: Summary Statistics

		Mean[i]	Standard deviation	Min	Max
		(1)	(2)	(3)	(4)
A.	**Contracts**				
(1)	Number of PDP plans per market	49	5.00	35	64
(2)	Number of insurers per market	16	0.60	14	17
(3)	Unweighted average PDP premium	$505	$59	$375	$643
B.	**Subsidies**				
(4)	CMS average national bid	$1,001	$45	$965	$1,060
(5)	CMS base consumer premium	$353	$26	$328	$383
(6)	CMS subsidy for average risk beneficiary	$648	$20	$631	$677
(7)	Low income benchmark threshold	$354	$24	$333	$388
C.	**Enrollment, millions[ii]**				
(8)	All Part D eligible	1.32	1.00	0.06	4.76
(9)	PDP enrollment, regular	0.26	0.18	0.01	0.70
(10)	PDP enrollment, low income[^]	0.24	0.20	0.01	1.02
(11)	MA-PD enrollment, regular	0.21	0.25	0.00	1.38
(12)	MA-PD enrollment, low income[^]	0.05	0.06	0.00	0.25
(13)	Employer sponsored coverage	0.20	0.16	0.01	0.48
(14)	Other coverage sources	0.17	0.11	0.01	0.48
(15)	No creditable coverage	0.19	0.13	0.01	0.58
D.	**Distribution of risk types among enrollees, %[iii]**				
(16)	Risk type 1, regular consumers	5%	2%	3%	15%
(17)	Risk type 2, regular consumers	20%	5%	11%	40%
(18)	Risk type 3, regular consumers	39%	5%	19%	52%
(19)	Risk type 4, regular consumers	7%	2%	2%	12%
(20)	Risk type 5, regular consumers	1%	0%	0%	1%
(21)	Consumers eligible for low-income subsidie	30%	6%	9%	45%

[i] Across 136 region-years (34 geographic markets in years 2007-2010)

[ii] Enrollment statistics based on CMS data on the whole market, not the analytic sample

[^] Counts of regular versus low income enrollment approximated using plan-level LIS counts enrollment

[iii] Based on the analytic sample as described in Section 4

Panel A based on CMS Part D Landscape files for years 2007 to 2010. A market is one of 34 Part D "regions" that cover 50 US States and the District of Columbia. "Insurers" are defined as Part D contracting organizations, which can have common ownership. Panel B is based on data from the annual releases by the Center for Medicare titled "Annual Release of Part D National Average Bid Amount and Other Part C & D Bid Information." Panels C and D report enrollment statistics across different types of Part D coverage (including not purchasing any coverage) as well as across consumer risk types within the primary types of Part D coverage considered in the paper - PDP and MA-PD.

Table 2: Results: Demand and Marginal Cost

	Beneficiaries not eligible for low-income subsidies					
	Risk type 1	Risk type 2	Risk type 3	Risk type 4	Risk type 5	LIS
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Parameters of the utility function						
Premium, $000						
Mean	3.16	3.25	2.46	2.56	2.56	-7.93
	(0.35)	(0.61)	(0.33)	(0.38)	(0.39)	(0.56)
Standard deviation	0.24	0.50	0.11	0.33	-0.07	-
	(0.50)	(0.39)	(0.62)	(0.67)	(0.72)	-
Standard deviation, inner option	0.67	6.92	-0.20	0.37	0.83	-
	(0.55)	(0.73)	(0.35)	(0.48)	(0.33)	-
Annual deductible, $000	-10.61	-11.36	-8.00	-7.97	-7.92	-
	(2.29)	(1.61)	(1.27)	(1.56)	(1.48)	-
Indicator for having any coverage in the gap						
Mean	3.83	3.75	1.82	3.38	4.62	-
	(2.08)	(1.32)	(1.12)	(0.86)	(1.14)	-
Standard deviation	2.05	-0.33	1.97	-0.19	0.87	-
	(0.85)	(0.98)	(0.53)	(0.65)	(0.72)	-
Number of most common drugs covered	47.87	34.03	40.85	43.92	60.00	1.05
	(17.57)	(11.11)	(14.11)	(10.85)	(8.68)	(7.39)
Measure of pharmacy network breadth	0.36	0.35	0.24	0.31	0.28	0.10
	(0.02)	(0.03)	(0.01)	(0.10)	(0.04)	(0.07)
Number of years the plan is on the market	1.09	1.24	1.07	1.05	0.92	0.48
	(0.19)	(0.12)	(0.10)	(0.13)	(0.12)	(0.04)
F-statistic first stage across all risk types	245					23
Panel B: Marginal cost projection for non-LIS market						
Annual deductible, $000	-0.19	-0.23	-0.29	-0.42	-0.52	
	(0.04)	(0.05)	(0.06)	(0.09)	(0.11)	
Number of most common drugs covered	-0.78	-0.30	-0.11	-2.11	-2.83	
	(2.21)	(2.69)	(3.42)	(4.95)	(6.17)	
Indicator for having any coverage in the gap	0.30	0.36	0.46	0.68	0.83	
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)	
Measure of pharmacy network breadth	-0.098	-0.14	-0.18	-0.28	-0.34	
	(0.04)	(0.05)	(0.07)	(0.10)	(0.12)	
Number of years the plan is on the market	0.03	0.04	0.04	0.06	0.08	
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	
Mean dependent variable (inverted MC, $000)	0.67	0.80	1.00	1.46	1.81	
Standard deviation dependent variable	0.16	0.20	0.25	0.36	0.44	
R-squared	0.77	0.76	0.76	0.76	0.76	
N	756	756	756	756	756	

The table reports parameter estimates for demand models as described in Section 3.1 (Panel A). The pa-rameters are reported separately for risk types 1 to 5 of regular consumers (columns 1-5) and consumerseligible for low income subsidies (column 6). We report standard errors in parentheses. All models include,but do not report: a constant, fixed effects for parent organizations, year fixed effects, geographic marketfixed effects,a dummy for an enhanced plan, the number of drugs on the plan's formulary, and the number ofdrugs placed in Tiers 1-2 of the formulary. F-statistics are reported for the first stage regressions of regular (column 1) and LIS (column 6) premiums on price instruments as described in Section 4. Panel B reportsthe results of a hedonic regression of marginal costs - for plans not distorted by LIS random assignmentincentives - estimated via the inversion of the first order conditions on plan characteristics in 2010. Theregression includes the same plan characteristics as the demand model for regular consumers.

Table 3: Results: Counterfactual Subsidy Mechanisms with Fixed Outside Option

		Observed Allocation		Remove cross- market links		Change market power		Subsidies proportional to bids			Flat voucher subsidies		
		Regular enrollees	LIS enrollees	No LIS link	No LIS, no MA-PD link	Independent Plans	Monopoly Ownership	p=5% of bid	p=32% of bid	p=95% of bid	$0	Optimal voucher: $800	$1,500
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
(1)	Consumer surplus, $M	2,298	2,642	2,678	3,028	3,080	2,443	11,032	3,613	974	950	2,955	10,742
(2)	Insurer profit, $M	559	-	1,062	1,205	1,154	1,923	35,260	3,821	40	27	1,311	2,605
(3)	**Consumer and producer surplus, $M**	2,857	2,642	3,739	4,233	4,235	4,367	46,293	7,434	1,014	977	4,266	13,347
(4)	Subsidy spending in PDP, $M	4,181	14,210	5,686	6,881	6,964	5,885	59,216	11,966	(35)	(56)	6,676	22,094
(5)	Reinsurance spending in PDP , $M	1,264	26,502	1,502	1,692	1,764	1,307	4,551	2,552	52	34	1,707	3,023
(6)	**Total government spending, $M**	5,445	40,712	7,188	8,573	8,728	7,192	63,768	14,518	17	(22)	8,383	25,117
(7)	Counterfactual subsidy spending if enrolled in MA-PD, $M	4,686	16,469	5,680	6,466	6,548	5,194	10,578	6,880	144	97	6,316	10,739
(8)	Counterfactual reinsurance spending if enrolled in MA-PD , $M	1,209	17,496	1,455	1,649	1,669	1,339	2,636	1,744	39	26	1,612	2,674
(9)	**Total opportunity cost of government spending, $M**	5,894	33,965	7,135	8,115	8,217	6,533	13,214	8,623	183	123	7,927	13,413
(10)	Total surplus; not accounting for opportunity cost of gov. spending, $M	(4,222)	(50,283)	(5,605)	(6,912)	(7,112)	(4,983)	(36,605)	(11,439)	992	1,005	(6,632)	(19,305)
(11)	Total surplus; accounting for opportunity cost of gov. spending, $M	3,441	(6,129)	3,671	3,638	3,570	3,510	(19,426)	(229)	1,230	1,165	3,674	(1,868)
(12)	Return on nominal dollar of gov. spending, $, no DWL of tax	(0.48)	(0.94)	(0.48)	(0.51)	(0.51)	(0.39)	(0.27)	(0.49)	59.13	(46.29)	(0.49)	(0.47)
(13)	Return on nominal dollar of gov. spending, $, with DWL of tax	(0.60)	(0.95)	(0.60)	(0.62)	(0.63)	(0.53)	(0.44)	(0.61)	45.26	(35.84)	(0.61)	(0.59)
(14)	Opportunity cost adjusted return on dollar of gov. spending, $	0.63	(0.15)	0.51	0.42	0.41	0.49	(0.30)	(0.02)	72.93	(54.00)	0.44	(0.07)
(15)	**Characteristics of the allocation**												
(16)	Inside option enrollment, '000	7,798	7,700	9,745	11,284	11,425	8,955	19,879	11,856	248	182	10,985	20,450
(17)	Inside option enrollment, percent of total market	36	79	45	52	53	42	92	55	1	1	51	95
(18)	Share of inside option enrollment by Risk Group 1 consumers, percent	6	-	6	7	7	7	7	6	0	0	7	7
(19)	Risk Group 2	17	-	18	18	17	19	21	15	67	72	18	23
(20)	Risk Group 3	59	-	60	61	61	58	62	64	2	1	61	60
(21)	Risk Group 4	16	-	14	13	13	15	10	13	31	27	14	9
(22)	Risk Group 5	1	-	1	1	1	1	1	1	0	0	1	1
(23)	Average weighted premium, $	510	25	474	438	441	450	201	603	1,448	1,531	444	78
(24)	Average weighted bid, $	1,129	1,051	1,170	1,174	1,175	1,254	4,011	1,885	1,524	1,531	1,244	1,473

Table reports the levels of consumer surplus, producer surplus, government spending, and total welfare under the observed allocation (columns 1 and 2) and under counterfactual allocations with a fixed outside option (columns 3 to 12). We compute these objects using estimates of demand and marginal costs (for columns 1 and 2), as well as simulations of counterfactual equilibria (columns 3 to 12). All quantities are computed as discussed in Section 3 and Appendix Section E. These baseline results assume that the cost of public funds (λ) is equal to 1.3. Negative quantities reported in parentheses.

Table 4: Results: Allocations and Welfare under Non-Market Mechanisms

		Planner	**Public Option**	
		Social Planner (1)	with subsidy (2)	without subsidy (3)
(1)	Consumer surplus, $M	3,258	3,005	970
(2)	Insurer profit, $M	(7,088)	-	-
(3)	**Consumer and producer surplus, $M**	(3,830)	3,005	970
(4)	Subsidy spending in PDP, $M	-	6,860	-
(5)	Reinsurance spending in PDP , $M	-	-	-
(6)	Additional subsidy spending in MA-PD, $M	-	-	-
(7)	**Total government spending, $M**	-	6,860	-
(8)	Counterfactual subsidy spending if enrolled in MA-PD, $M	6,943	5,599	154
(9)	Counterfactual reinsurance spending if enrolled in MA-PD , $M	1,770	1,433	41
(10)	**Total opportunity cost of government spending, $M**	8,713	7,032	195
(11)	Total surplus; not accounting for opportunity cost of gov. spendin	(3,830)	(5,912)	970
(12)	Total surplus; accounting for opportunity cost of gov. spending, $I	5,371	3,229	1,223
(13)	Return on nominal dollar of gov. spending, $, no DWL of tax	-	(0.56)	-
(14)	Return on nominal dollar of gov. spending, $, with DWL of tax	-	(0.66)	-
(15)	Opportunity cost adjusted return on dollar of gov. spending, $	-	0.47	-
(16)	**Characteristics of the allocation**			
(17)	Inside option enrollment, '000	12,374	10,237	322
(18)	Inside option enrollment, percent of total market	57	47	1
(19)	Share of inside option enrollment by Risk Group 1 consumers, per	7	9	1
(20)	Risk Group 2	18	19	71
(21)	Risk Group 3	61	59	10
(22)	Risk Group 4	13	12	19
(23)	Risk Group 5	1	1	0
(24)	Average weighted premium, $	374	87	728
(25)	Average weighted bid, $	374	758	728

Table reports the level of consumer surplus, producer surplus, government spending, and total welfare under counterfactual allocations without market mechanisms. The non-market mechanisms are defined in Section 5.1. To compute these objects, we use estimates of demand, marginal costs, and the derivation of the social planner's problem in Appendix Section E. All quantities are computed as discussed in Section 3 and Appendix Section E. These baseline results assume that the cost of public funds (λ) is equal to 1.3. Negative quantities are reported in parentheses.

Table 5: Results: Counterfactual Subsidy Mechanisms with Adjusted Outside Option

		Remove cross-market links		Change market power		Flat voucher subsidies		
		No LIS link	No LIS, no MA-PD link	Independent Plans	Monopoly Ownership	$0	Optimal voucher: $1,200	$1,500
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1)	Consumer surplus, $M	3,504	4,189	4,227	4,410	(12,460)	9,577	11,327
(2)	Insurer profit, $M	873	914	881	1,220	832	1,992	2,173
(3)	**Consumer and producer surplus, $M**	4,378	5,103	5,108	5,630	(11,628)	11,569	13,500
(4)	Subsidy spending in PDP, $M	4,725	5,269	5,366	4,266	(920)	15,596	19,707
(5)	Reinsurance spending in PDP , $M	1,252	1,306	1,368	957	1,152	2,335	2,831
(6)	Additional subsidy spending in MA-PD, $M	442	802	789	1,415	(10,344)	943	674
(7)	**Total government spending, $M**	6,419	7,377	7,522	6,638	(10,112)	18,874	23,212
(8)	Counterfactual subsidy spending if enrolled in MA-PD, $M	5,075	5,545	5,649	4,363	(448)	12,297	13,493
(9)	Counterfactual reinsurance spending if enrolled in MA-PD , $M	1,203	1,257	1,280	962	1,075	2,274	2,456
(10)	**Total opportunity cost of government spending, $M**	6,278	6,802	6,929	5,325	627	14,571	15,949
(11)	Total surplus; not accounting for opportunity cost of gov. spending, $M	(3,967)	(4,487)	(4,671)	(3,000)	1,517	(12,967)	(16,675)
(12)	Total surplus; accounting for opportunity cost of gov. spending, $M	4,194	4,356	4,337	3,922	2,332	5,975	4,058
(13)	Return on nominal dollar of gov. spending, $, no DWL of tax	(0.32)	(0.31)	(0.32)	(0.15)	0.15	(0.39)	(0.42)
(14)	Return on nominal dollar of gov. spending, $, with DWL of tax	(0.48)	(0.47)	(0.48)	(0.35)	(0.12)	(0.53)	(0.55)
(15)	Opportunity cost adjusted return on dollar of gov. spending, $	0.65	0.59	0.58	0.59	(0.23)	0.32	0.17
(16)	**Characteristics of the allocation**							
(17)	Inside option enrollment, '000	7,769	8,195	8,355	6,016	6,785	16,231	17,774
(18)	Inside option enrollment, percent of total market	36	38	39	28	31	75	82
(19)	Share of inside option enrollment by Risk Group 1 consumers, percent	6	6	6	6	6	7	7
(20)	Risk Group 2	17	17	17	18	17	16	16
(21)	Risk Group 3	59	60	60	56	58	64	65
(22)	Risk Group 4	16	16	16	19	18	11	11
(23)	Risk Group 5	1	1	1	1	1	1	1
(24)	Average weighted premium, $	466	427	429	414	1,224	94	86
(25)	Average weighted bid, $	1,162	1,164	1,164	1,219	1,224	1,286	1,461
(26)	Outside option adjustment, $	56	84	84	113	(677)	203	218

Table reports the level of consumer surplus, producer surplus, government spending, and total welfare under counterfactual allocations with an endogenously adjusted outside option. To compute these objects, we use estimates of demand, marginal costs, and simulations of counterfactual equilibria. All quantities are computed as discussed in Section 3 and Appendix Sections E and F. These baseline results assume that the cost of public funds (λ) is equal to 1.3. Negative quantities reported in parentheses.

APPENDIX

A Institutional Context and Data

Institutional Context Congress expanded Medicare to include prescription drug coverage via Medicare Part D in 2006. 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.

The supply-side of the Part D program has a unique, and controversial, design. Oliver, Lee and Lipton (2004) discuss the political origins of Part D and its mixed reception in the first years of the program, particularly among consumers. 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 A1; 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.[25] 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 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. There are several nuances buried in the set-up of the bidding procedure that are important for insurers' incentives and will

[25] See Table IV.B11 of 2012 Trustees of Medicare Annual Report.

Figure A1: Standard Defined Benefit in Medicare Part D

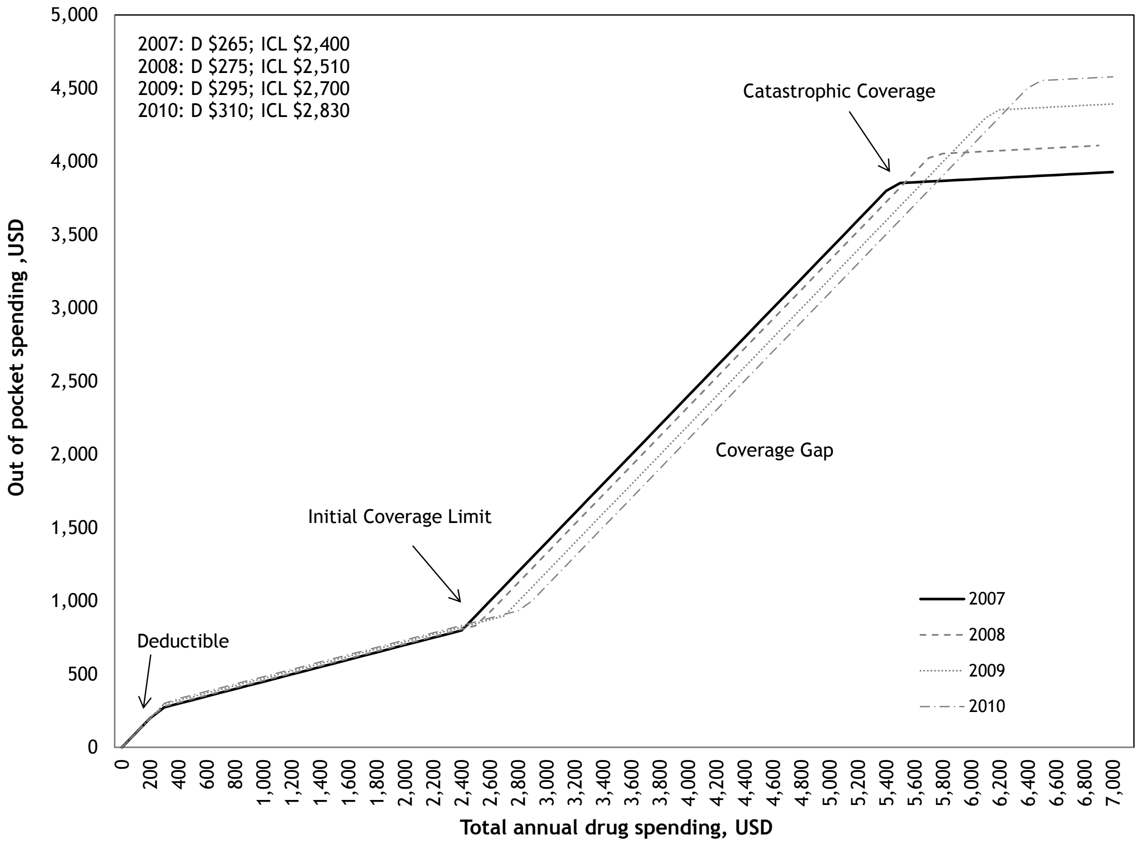

Insurers in the Medicare Part D program are required to provide coverage that gives at least the same actuarial value as the Standard Defined Benefit (SDB) that is illustrated in this Figure for years 2007-2010. The SDB design features a deductible, a co-insurance rate of 25 percent up to the initial 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 (on the x-axis) into the out-of-pocket expenditure (on the y-axis). The legend in the top left corner specifies the key parameter values of the contract and their evolution over time: deductible level ("D") and the start of the coverage gap phase ("ICL").

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. 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. 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 enrollment-weighted 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. CMS payments to insurers are scaled by this risk score to reflect higher expected expenditures that insurers would incur for sicker enrollees. For a consumer with risk score r_i, CMS pays the insurer the bid times the risk score, minus consumer premium. While the premiums do not vary across consumers, CMS payment is higher for less healthy 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 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.

MA-PD plans participate in the bidding mechanisms and conceptually are paid the same subsidy. In practice, MA-PD plans can use the subsidy payments from the medical part of MA to reduce consumer premiums for MA-PD plans, so that consumer-facing premiums are often zero for the pharmaceutical part of the MA program (Curto et al., 2015; Starc and Town, 2015).

Data and Sample Our primary data source for analysis are Research Identifiable Files containing a 20 percent random sample of all pharmaceutical claims for Medicare beneficiaries from years 2007-2010. The dataset (Master Beneficiary Summery File) contains information about age, sex, place of residence, health conditions, and Part D enrollment for each beneficiary in the 20 percent sample. For those individuals that have purchased some type of Part D coverage, we observe detailed information about pharmaceutical claims (Part D Drug Event File). For each claim, we observe the date of purchase, the pharmacy, the prescribing physician, the total cost of transaction at the pharmacy, how much of the claim was paid by the consumer, by the plan, and by CMS. We can further observe if the consumer was eligible for low income subsidies. These data are described in detail by CMS Research Data Assistance Center (ResDAC): www.resdac.org.

We restrict the 20 percent sample as follows. We keep individuals that reside in 50 US states. We further keep individuals that either purchased a stand-alone Medicare Part D PDP plan, or a Medicare Advantage MA-PD plan, or no Part D coverage. This restriction excludes individuals who were eligible for other types of pharmaceutical coverage - primarily through employer-sponsored Part D plans. This set of individuals constitutes our analytic sample.

We enrich these data with publicly available information on the Part D plan options that were available to each consumer based on their geographic market. For each plan option we observe a vector of detailed characteristics, including deductibles, coverage in the gap, and insurer brands, as well as consumer-facing premiums. These CMS Medicare Part D Landscape files for years 2007-2010 are provided directly by CMS at www.cms.gov.

Last, we use public information on total and plan-level Part D enrollment as reported on www.cms.gov to validate the estimates based on our 20 percent sample. The same source of information provides data on average national bids, consumer premiums, and low income benchmark thresholds. We use these pieces of information to infer bids for each Part D plan, as well as which plans were eligible for LIS random assignment.

For details on data sources, file names and processing programs, please see the package with replication programs that accompanies the paper.

B Consumer Risk Types

Our supply side model allows for a plan's marginal cost to change as a function of the plan's risk pool; in other words, the supply-side model allows for adverse or advantageous selection. To facilitate differential risk sorting across plans, our demand model allows for consumer preferences to vary with their health risk. To operationalize endogenous marginal costs as an equilibrium outcome, we start by constructing a continuous one-dimensional measure of each consumer's risk. The risk measure that we construct is similar in spirit to the risk score that CMS uses for risk-adjustment. We then discretize the risk space to enable computational feasibility.

We start by generating a continuous measure of predictable risk for each Medicare beneficiary in the data. For all consumers on the market that are included in our 20 percent sample, we observe information about their health status, measured by the indicators for the presence of 66 chronic conditions. To map the indicators for chronic conditions and basic demographic information into a one-dimensional pharmaceutical spending risk score, we estimate a linear relationship between total pharmaceutical spending and information about individual's chronic conditions, age, sex, race, and eligibility for low income subsidies. The model is estimated using observations on individuals who had Part D coverage and hence their pharmaceutical claims were recorded in the data. We then use the estimated regression coefficients to predict total pharmaceutical spending among all potential Part D consumers. For each individual, we construct a risk score measure that is equal to the ratio of the individual's predicted spending to the average predicted spending in the sample. By construction, the average consumer receives a risk score of 1.

In the next step we discretize the constructed risk scores. We divide all potential Part D enrollees into six discrete risk groups. As the distribution of risk scores has thin tails and is concentrated around the mean, we define the first and fifth risk score groups to be the bottom 5th percentile and top 5th percentile of risk scores among beneficiaries that are not eligible for low income subsidies. The second risk group are individuals with risk scores between 5th and 25th percentile; third risk group - 25th to 75th percentile, and fourth risk group - 75th to 95th percentile. We let all LIS-eligible beneficiaries be in a separate, sixth, risk group. We cannot observe individual's LIS eligiblility if the individual is not enrolled in a Part D plan. Hence, we assume that all consumers that are not observed having a Part D plan are not eligible for LIS subsidies. This appears to be a reasonable assumption, as LIS eligiblity for many individuals is determined automatically and those who are not actively choosing a Part D plan get randomly assigned to one.

For each of the six (including LIS) risk groups, we define two objects of interest. First, we compute the average risk score in each group. The resulting averages are: 0.17 risk score for the first risk group, 0.47 for the second, 0.81 for the third, 1.63 for the fourth, 2.49 for the fifth, and 1.78 for the LIS. In the process of risk-adjustment, CMS constructs similar risk scores for each individual. These "real" risk scores are then used to multiply insurers' bids before CMS determines its payments to the insurer. CMS premium subsidy is equal to the difference between the insurer bid multiplied by a risk score net of the beneficiary premium. To replicate this idea and to allow insurers to collect higher revenue when enrolling higher cost individuals, we use the average risk scores per risk type as reported above to multiply insurer bids and compute insurer revenue.

In the next step, we define how expected costs vary across five risk groups. Since we observe only one equilibrium price per plan, we have to make parametric assumptions on how costs across six risk group types relate to each other, as there is only one unknown cost parameter that is identified in the marginal cost inversion system. We assume that risk-type specific costs are related multiplicatively to each other. We normalize the marginal cost for risk group one as the baseline marginal cost c. We then assume that the marginal cost for risk groups two to six (the LIS) are equal to $\kappa_2 c$ to $\kappa_6 c$. We estimate κ's from the claims data. For each risk type, we compute the average *plan* payment, including payments by PDP and MA-PD type plans. The insurance plan payment per person is the residual of the total annual individual spending on drugs net of patient out of pocket payments for cost-sharing, net of reinsurance payments for especially high risks by the government, and net of government cost-sharing subsidies for the LIS-eligible beneficiaries. We normalize the resulting average payment in all risk types by the payment in risk type 1. The resulting estimates of κ are: 1 for risk type 1, 1.19 for risk type 2, 1.50 for risk type 3, 2.18 for risk type 4, 2.70 for risk type 5, 1.93 for the LIS-eligible beneficiaries. Applying these estimated multipliers allows us to reduce the marginal cost inversion problem to a one equation in one unknown, while at the same time endogenizing the marginal costs of Part D plans to equilibrium allocation of risks across plans.

C Consumer Inertia

The literature has documented consumer inertia in the choice of Medicare Part D plans (Ericson, 2014; Ho et al., 2015; Polyakova, 2016; Wu, 2016; Heiss et al., 2016). Consumers tend to choose their plan when entering the Part D program for the first time, and then only infrequently make changes to their plans. To account for inertia in demand, we take a reduced form approach and include 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. As we illustrate in what follows, our reduced-form specification corresponds to an explicit structural model of inattention and choice.

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 j depends on its own share from the previous period as follows:

$$\hat{s}_{j,t}(p, s_{j,t-1}) = \alpha M Pr_{j,t}(p) + (1-\alpha)s_{j,t-1}, \tag{7}$$

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 it 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 7 with respect to its premium only has the current set of active choosers in it:

$$\frac{\partial \hat{s}_{j,t}(p, s_{j,t-1})}{\partial p_{j,t}} = \alpha M \frac{\partial Pr_{j,t}(p)}{\partial p_{j,t}}, \tag{8}$$

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 8 as:

$$\frac{\partial \hat{s}_{j,t}(p, s_{j,t-1})}{\partial p_{j,t}} = \alpha(T) M \frac{\partial Pr_{j,t}(p)}{\partial p_{j,t}} = M \frac{\partial Pr_{j,t}(p, \beta(T))}{\partial p_{j,t}}, \tag{9}$$

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.

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 have re-estimated our key counterfactual results for a 20 percent radius around our marginal cost estimates. Our qualitative conclusions are not sensitive to this specification check.

D Profit Function

This section provides a detailed description of how we arrive at the profit function in Equation 3 of the paper. We start with a description of the flow of payments in Part D and set up a general profit function that can incorporate a variety of regulatory interventions in this market. 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. The premium does not vary across consumers and is determined as follows. CMS takes a (lagged) enrollment-weighted average of all bids submitted by all PDP and MA-PD plans across the country. It then declares a pre-specified share of this average for the given year (for example, 36 percent in year 2010) as the base consumer premium. The actual premium is then equal to the base premium plus the difference between the plan's bid and the average national bid.

The consumer premium is augmented with an individual-plan-specific subsidy, z_{ij}, from the government. This subsidy is equal to the bid multiplied by a measure of the enrollee's ex-ante health risk - r_i) - net of consumer premium. For an average-risk beneficiary with $r_i = 1$, the sum of the premium and government subsidy is equal to the bid that the firm submitted for that plan. For an individual with a risk score above or below the average, the insurer collects $r_i * b_j$ - out of this amount, p_j is paid by the consumer, and $z_{ij} = r_i * b_j - p_j$ is paid as the government subsidy. Since consumer premium depends on the average bid among all PDP and MA-PD plans, $\bar{b}$, we can write the subsidy as a function of the bid, the average bid, and individual-specific health risk: $z_{ij}(b_j, \bar{b}, r_i)$. The individual-level risk adjustment in the subsidy is intended to make all consumers look equally profitable to firms in order to reduce incentives for risk-based selection.

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 reinsurance subsidies. 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 happen at the end of the year, and restrict the downside (but also upside) risk of enrolling extremely costly individuals for the insurers. Medicare Part D Manual provides more details. 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).) 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:

$$\pi_j(b_j; b_{-j}) = \Gamma\left[\sum_{i\in j}\left(p_j(\bar{b}, b_j) + z_{ij}(b_j, \bar{b}, r_i) - c_{ij}(r_i, \phi_j)\right)\right], \tag{10}$$

where the summation is taken over all individuals enrolling in the plan.

As the sum of the premium and the subsidy is by construction equal to the risk-adjusted bid submitted by insurer to Medicare, $p_j(\bar{b}, b_j) + z_{ij}(b_j, \bar{b}, r_i) = r_i * b_j$, we can re-write the ex ante expected profit of plan j for all consumers with risk level r as:

$$\pi_j^r(b_j; b_{-j}) = M_r s_{rj}(b)(r b_j - c_{rj}), \tag{11}$$

where $s_j^r(b)$ is the market share of plan j among consumers of risk r and M is the market size. We emphasize that s_j^r encapsulates all of the regulatory details involved in turning bids, b, into plan-specific market shares. To operationalize the analysis, we discretize the risk type space into five risk types among regular enrollees and LIS enrollees as a separate risk type. Let t index risk types of regular enrollees. Let θ_t denote the average risk score among type t enrollees. We can then re-write the risk-type level profit function as:

$$\pi_j^t(b_j; b_{-j}) = M_t s_{jt}(b)(\theta_t b_j - c_{jt}), \tag{12}$$

We now expand this expression to allow for multi-plan insurance organizations that offer plans to all risk types on the market, including the LIS consumers. The structure of profit from LIS enrollment is specified as entirely symmetric to the regular enrollees. We denote quantities related to regular enrollees of risk type t with superscript R, and quantities related to the LIS part of the market with superscript LIS. The profit function for insurer f offering a portfolio of $j \in J_f$ plans across all consumer types becomes:

$$\pi_f(b) = \sum_{j \in J_f} \left(\sum_{t=1}^{5} \left[M_t^R s_{jt}^R(b)(\theta_t^R b_j - c_{jt}^R) \right] + M^{LIS} s_j^{LIS}(b)(\theta^{LIS} b_j - c_j^{LIS}) \right) \tag{13}$$

Firms maximize profits by choosing bid b for each insurance plan j they offer.

Equation 13 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_j^R, which is not set directly by insurers, but rather depends on the bids of other insurers in a non-linear fashion:

$$p_j^R = \max\left\{0, b_j - \bar{b} + \zeta \bar{b}\right\}, \tag{14}$$

where $\bar{b}$ is the enrollment-weighted average bid of all plans in the entire US and ζ is the share of the average bid allocated to baseline consumer premiums. 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 0.36. 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 consumer that make active choices of plans ("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.[26]

We make two assumptions to arrive at a tractable first order condition. First, we 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. Second, we assume that only insurers that are not competing for the randomly assigned LIS beneficiaries can be characterized as playing a Nash-Bertrand game. For these plans, the first-order condition is:

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

$$\frac{\partial \pi_f}{\partial b_j} = \sum_{t=1}^{5} \left[\theta_t M_t^R s_{jt}^R(b) + (\theta_t b_j - c_{jt}^R) M_t^R \frac{\partial s_{jt}^R(b)}{\partial b_j} + \sum_{k \neq j \in J_f} (b_k - c_{kt}^R) M_t^R \frac{\partial s_{kt}^R(b)}{\partial b_j} \right]$$

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

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 decision-making. The market size affects the relative effects on profit from enrolling beneficiaries of different risk types. The incentives are further complicated by the fact that we allow for adverse or advantageous selection and endogenous marginal costs. The profitability for isnruers varies across risk types, since we do not impose that additional revenues from CMS for higher risk enrollees through the risk adjustment program offset the cost differences. To close the model, we assume that the expected marginal costs for risk types 2 to 5 and LIS are a first order polynomial with an intercept equal to zero relative to risk type 1. We denote the slope parameters with κ_t and κ_{LIS}. Collecting terms in vector notation, we can re-write the first order condition for an insurer offering plans not distorted by LIS incentives as follows:

$$\sum_{t=1}^{T} \left[\theta_t M_t^R s_t^R - \Omega_t^R (\theta_t b - \kappa_t c) \right] + \theta^{LIS} M^{LIS} s^{LIS} - \Omega^{LIS} (\theta^{LIS} b - \kappa^{LIS} c) = 0. \tag{15}$$

where

$$\Omega_{kjt}^R = \begin{cases} -M_t^R \frac{\partial s_{tj}^R(b)}{\partial b_k} & \text{if } \{j,k\} \in J_f, \\ 0 & \text{else,} \end{cases} \tag{16}$$

and

$$\Omega_{kj}^{LIS} = \begin{cases} -M^{LIS} \frac{\partial s_j^{LIS}(b)}{\partial b_k} & \text{if } \{j,k\} \in J_f, \\ 0 & \text{else.} \end{cases} \tag{17}$$

We use these first-order conditions to compute (by inversion for non-distorted plans and using a hedonic projection for distorted plans) for marginal costs for each plan j. These cost estimates are in turn used as inputs for computing the counterfactual equilibria.

E Welfare Function

For regular enrollees, total welfare in the Medicare Part D PDP market is comprised of three pieces: consumer surplus (CS), insurer profits (Π), and government spending (G):

$$W = CS + \Pi - \lambda G, \tag{18}$$

where λ is the social cost of raising revenues to cover government expenditures, G. All three pieces of the welfare function are calculated relative to the outside option. For consumer surplus the normalization to the outside option (buying an MA-PD plan or not purchasing Part D insurance) follows directly from the utility model. For producer surplus, the insurer pricing decision 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 outside of the Part D program. 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 the PDP and MA-PD parts of the market, we consider government spending on PDP net of what it would have spent on the same individual elsewhere. We conservatively assume that the outside option for the government is subsidizing the same consumers in the MA-PD market. 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 ω_i from plan characteristics, including the premium, takes the following form:

$$CS(\omega_i) = \frac{1}{\alpha_i}\left[\gamma + \ln\left[1 + \sum_{j=1}^{J} \exp(v_{ij}(\omega_i))\right]\right], \tag{19}$$

where γ is Euler's constant, and v_{ij} is the deterministic component of utility for person i from plan j (utility net of the idiosyncratic shock).[27] We integrate out over the unobserved taste heterogeneity to obtain consumer surplus for each consumer risk group t:

$$CS_t = \int CS(\omega_t) dF(\omega_t). \tag{20}$$

The second piece of the welfare calculation is producer surplus that we compute using

[27] Euler'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.

Equation 3.

The last piece of the welfare calculation is government spending. In our welfare calculations, we weigh the government spending with the shadow cost of public funds, commonly estimated at $\lambda = 1.3$.[28] We compute both the nominal government spending in the PDP program, as well as how much extra 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}). We allow government spending to vary by consumer risk type, which reflects differential risk adjustment and re-insurance payments across consumers of different health.

Adding the three parts of the welfare function back together, we have the following measure of total surplus for regular consumers (an analogous expression applies to LIS consumers):

$$\begin{aligned} W = & \overbrace{\sum_{t=1}^{t=5} M_t \int \frac{1}{\alpha_t} \left(\gamma + \ln \left[1 + \sum_{j=1}^{J} \exp(v_{jt}(\omega_t)) \right] \right) dF(\omega_t)}^{\text{Consumer Surplus (CS)}} + \\ & \overbrace{\sum_{j \in J} \left(\sum_{t=1}^{5} [M_t s_{jt}(b)(\theta_t b_j - \kappa_t c_j)] \right)}^{\text{Producer Profit } (\Pi)} - \\ & \overbrace{\lambda \left(\sum_{t=1}^{t=5} \left[\sum_{j=1}^{J} (G_{jt}^{PDP} - G_t^{MAPD}) s_{jt}(p) M_t \right] \right)}^{\text{Social Cost of Government Spending (G)}} . \end{aligned} \tag{21}$$

This welfare function describes the surplus for the private market, where firms administer insurance contracts. To arrive at the social planner's objective function, we include "profits" in the government spending computation, as under a social planner, prices are set by fiat and there are no private profits. Social planner's problem is then to maximize the following welfare function:

[28]See, for example, Hausman and Poterba (1987).

$$W^{SP}(p) = \overbrace{\sum_{t=1}^{t=5} M_t \int \frac{1}{\alpha_t} \left(\gamma + \ln \left[1 + \sum_{j=1}^{J} \exp(v_{jt}(\omega_t)) \right] \right) dF(\omega_t)}^{\text{Consumer Surplus (CS)}} +$$

$$\lambda \left[\overbrace{\sum_{j \in J} \left(\sum_{t=1}^{5} \left[M_t^R s_{jt}^R(p)(\theta_t^R p_j - \kappa_t^R c_j) \right] \right)}^{\text{Producer Profit }(\Pi)} - \overbrace{\left(\sum_{t=1}^{t=5} \left[\sum_{j=1}^{J} (G_{jt}^{PDP} - G_t^{MAPD}) s_{jt}(p) M_t \right] \right)}^{\text{Social Cost of Government Spending (G)}} \right] \tag{22}$$

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. The optimal price vector 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:

$$\frac{\partial CS(p)}{\partial p_j} = \sum_{t=1}^{t=5} \left[\int M_t \frac{1}{\alpha_t} \left[\frac{-\alpha_t \exp(v_{jt}(\omega_t))}{1 + \sum_{k=1}^{J} \exp(v_{kt}(\omega_t))} \right] dF(\omega_t) \right] = -\sum_{t=1}^{t=5} M_t s_{jt}(p) \tag{23}$$

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

$$\frac{\partial \Pi(p)}{\partial p_j} = \sum_{t=1}^{t=5} \left[\lambda M_t s_{jt}(p) \theta_t + \lambda M_t \sum_k (\theta_t p_k - \kappa_t c_k) \frac{\partial s_{kt}(p)}{\partial p_j} \right] \tag{24}$$

The derivative of government spending with respect to p_j is:

$$\frac{\partial GS(p)}{\partial p_j} = \sum_{t=1}^{t=5} \left(-\lambda \left[\sum_k (G_{kt}^{PDP} - G_t^{MAPD}) \frac{\partial s_{kt}(p)}{\partial p_j} M_t \right] \right) = \tag{25}$$

$$\sum_{t=1}^{t=5} \left(-\lambda \left[\sum_k \Delta G_{kt} \frac{\partial s_{kt}(p)}{\partial p_j} M_t \right] \right) \tag{26}$$

Summing these terms, we obtain:

$$\frac{\partial W^{SP}(p)}{\partial p_j} = \sum_{t=1}^{t=5} M_t \left((\lambda \theta_t - 1) s_{jt}(p) + \lambda \sum_k (\theta_t p_k - \kappa_t c_k - \Delta G_{kt}) \frac{\partial s_{kt}(p)}{\partial p_j} \right) \tag{27}$$

A decrease in consumer surplus in response to an increased price ($-s_{jt}(p)$) is offset, up to the cost of transferring public funds, by an increase in profit in the product market ($\lambda\theta_t s_{jt}(p)$). The degree of offset varies by consumer risk type.

The first-order conditions can be simplified using vector notation:

$$\sum_{t=1}^{t=5} M_t \left[(\lambda\theta_t - 1)s_t(p) + \lambda\Omega_t(p)(\theta_t p - \kappa_t c - \Delta G_t)\right] = 0 \tag{28}$$

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

$$\Omega_{ijt}(p) = \frac{\partial s_j t(p)}{\partial p_i}. \tag{29}$$

It follows that optimal prices for the social planner's case are given by the following:

$$p^{SocialPlanner} = \left(\sum_{t=1}^{t=5} \lambda M_t \Omega_t(p)\theta_t\right)^{-1} \left(\sum_{t=1}^{t=5} M_t \left[(1 - \lambda\theta_t)s_t(p) + \lambda\Omega_t(p)(\kappa_t c + \Delta G_t)\right]\right) \tag{30}$$

Price is set to balance the inside and outside option enrollment for each consumer risk type. In particular, the size of each risk type market as well as the difference in risk adjustment (θ_t) versus cost factors across risk types (κ_t) play a central role in determining the social planner's allocation.

Cost of public funds A key parameter in our calculations in the social cost of government funds, which is set to 1.3 in our baseline analysis. In Table E1, we report the estimates of total surplus (accounting for the opportunity cost of government funds) for each counterfactual for $\lambda \in \{1, 1.7, 2\}$. Increasing the cost of public funds has the general effect of decreasing overall welfare in most cases. The optimal voucher shifts up by \$100 to \$900 at $\lambda = 1$ in the fixed outside option case, and remains the same at \$1,200 in the adjusted outside option case. At $\lambda = 2$, the optimal voucher shifts down by \$100 to \$700 in the fixed outside option case. In the case of adjusted outside option, making government payments more socially costly leads to a knife-edge case, where the optimal strategy is not to subsidize the market at all, since there is not enough willingness to pay for the PDP program, so that any subsidy is costlier than the utility loss from not subsidizing. The latter remains the case as long as the public cost of government funds is above 1.5 per 1 dollar of government spending.

Table E1: Sensitivity to the Cost of Public Funds Parameter

	Counterfactuals	λ=1	λ=1.3	λ=1.7	λ=2
	(1)	(2)	(3)	(4)	(5)
Panel A	**Fixed outside option**				
(1)	Observed allocation, $M	3,306	3,441	3,620	3,755
(2)	No LIS link, $M	3,687	3,671	3,650	3,634
(3)	No LIS , no MA-PD link, $M	3,775	3,638	3,455	3,317
(4)	Independent plans, $M	3,724	3,570	3,366	3,212
(5)	Monopoly ownership, $M	3,708	3,510	3,247	3,049
(6)	$0 voucher, $M	1,121	1,165	1,222	1,266
(7)	Optimal voucher, $M	3,891	3,674	3,680	3,771
	Optimal voucher level, $	*900*	*800*	*700*	*700*
(8)	$ 1,500 voucher, $M	1,643	-1,868	-6,549	-10,060
Panel B	**Adjusted outside option**				
(9)	No LIS link, $M	4,237	4,194	4,138	4,096
(10)	No LIS , no MA-PD link, $M	4,529	4,356	4,127	3,954
(11)	Independent plans, $M	4,515	4,337	4,100	3,922
(12)	Monopoly ownership, $M	4,316	3,922	3,396	3,002
(13)	$0 voucher, $M	-889	2,332	6,628	9,850
(14)	Optimal voucher, $M	7,266	5,975	6,628	9,850
	Optimal voucher level, $	*1,200*	*1,200*	*0*	*0*
(15)	$ 1,500 voucher, $M	6,237	4,058	1,153	-1,025
Panel C	**Non-market mechanisms**				
(16)	Public option with subsidy, $M	3,178	3,229	3,298	3,350
(17)	Public option without subsidy, $M	1,165	1,223	1,301	1,360

Table reports the level of welfare (accounting for the opportunity cost of government funds) under counterfactual allocations with a fixed outside option, endogenously adjusted outside option, and non-market mechansims for different levels of the cost of public fund parameter (λ). Welfare is computed exactly as in baseline counterfactuals, only varying the cost of public funds parameter. The baseline results assume that the cost of public funds (λ) is equal to 1.3 - these results are replicated in column (2).

F Outside Option Adjustment

We proceed in several steps to calculate the adjustments in the value of MA-PD coverage for consumers that allow for changes in the outside option in counterfactual equilibria. There are slight differences across different counterfactuals, so we describe them separately.

We start with counterfactuals that 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 first compute the counterfactual premium subsidy for PDP plans. This subsidy 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 using subsidization formulas that link bids and premiums. 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 the degree of "buy down" by plan is not publicly available. To circumvent this data gap, 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 bid. In addition, we assume that MA-PD plans would apply the same "buy-down" on top of any counterfactual 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. We proceed analogously in counterfactuals where subsidies are set as flat vouchers. The key 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 imputed MA-PD bids. We then compute the resulting difference in average MA-PD premiums for consumers, which is the outside option adjustment recorded in row (26) of Table 5. We use this adjustment as our measure of the change in the value of the outside option. In practice, the increase (decrease) in the attractiveness of the outside option is computationally implemented as a symmetric decrease (increase) in the value of the inside option (the constant for the inside option choice is adjusted).

G Subsidy Design in LIS Market

Our baseline set of counterfactuals removes the LIS enrollees from the market. As discussed in the paper, this was done for both pedagogical and technical reasons. Naturally, however, determining optimal ways to provide subsidies to both markets is an important question to consider. One must confront two questions when thinking about joint subsidy setting for LIS and regular beneficiaries. First, what is 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, requiring that insurers submit only one bid for both markets, combined with the LIS eligibility threshold, serves as a brake on premiums. It introduces an elasticity of demand 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. Earlier work in Decarolis (2015) explores these issues in detail.

To shed light on the counterfactual mechanisms that could maximize joint surplus of LIS and regular markets, we reintroduce LIS enrollees into the regular market in a supplementary set of counterfactuals. We 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 voucher subsidies for LIS and regular consumers. We ran the analysis on a matrix of LIS and regular vouchers that range from \$0 to \$1000 in \$100 increments. We find that under the double voucher counterfactuals, the optimal voucher for regular consumers is \$800, which is the same as the optimal voucher when just considering the regular market; for the LIS consumers, lower vouchers result in general lead to higher surplus due to the very high government cost of subsidizing this market. The differences across total surplus along the LIS voucher dimension are relatively small as long as the LIS voucher is lower than $800. It follows that setting the LIS voucher to be equal to the optimal voucher for regulars does not lead to a significant decrease in total surplus.

H 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. Taking this average bid, $\bar{b}$, as fixed, we then solve for the vector of first-order conditions. 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 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 for each risk type of regular enrollees was standard with the exception of imposing the lognormality of the price coefficient. All programs and instructions on obtaining data are publicly available in the technical supplement.