Information Frictions and Insurer Plan Design: Evidence from Medicare Advantage

Raymond Kluender, Evan Mast*

First Version: December 2015

This Version: June 2017

Abstract

Information frictions in markets for insurance affect not only the choices consumers make, but also the menu of plans insurers offer. We illustrate this observation using an information friction in Medicare Advantage—beneficiaries pay two premiums, and one is much more salient. We begin by estimating demand and finding a larger elasticity for the salient versus non-salient premium. Next, we show that a model of insurer plan design produces simulated premiums matching the observed distribution when accounting for differential salience, but not when assuming equal elasticities across the two premiums. Finally, we simulate how plan enrollment would change if the friction were removed. Consumer surplus increases by $73/year when allowing insurers to redesign their plans, versus only $5/year holding supply fixed.

JEL Codes: I11, I13, D12, G22

Keywords: Medicare, Health Insurance Exchanges, Salience, Medicare Advantage, ACA

*Kluender: Depart of Economics, MIT, kluender@mit.edu; Mast: Department of Economics, Stanford University, evanemast@gmail.com. Thanks to Jay Bhattacharya, Tim Bresnahan, Mark Duggan, Liran Einav, Amy Finkelstein, Jonathan Gruber, Vilsa Curto, Atul Gupta, Jonathan Petkun, Pietro Tebaldi, and Dan Waldinger. Mast gratefully acknowledges support from the Ric Weiland Graduate Fellowship and the E.S Shaw and B.F. Haley Fellowship. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 (Kluender). Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation

1 Introduction

Tens of millions of Americans purchase health insurance via online exchanges that offer coverage through Medicare Advantage, Medicare Part D, and the Affordable Care Act. A large body of literature suggests that individuals find these exchanges confusing and struggle to fully understand the characteristics of the available plans. This confusion may affect not only the plans that consumers choose, but also the set of plans available. Assuming insurers design their plans to maximize profit given consumer behavior—whether or not that behavior is optimal—information frictions should affect the plan attributes insurers choose.

We study an information friction on the Medicare Advantage (MA) exchange that appears to have a large effect on insurer plan design. MA Beneficiaries must pay two premiums (Part B and Part C Supplemental), and the Part C supplemental premium is much more salient than the Part B premium. Our primary motivating fact, previously noted by Stockley et al. (2014) and Newhouse & McGuire (2014), is shown in Figure 1, which displays the distribution of total premiums in Medicare Advantage in 2011. Consistent with a significant demand response to the salient but not the nonsalient premium, there is a massive peak at $96.40, where plans are charging the maximum Part B premium and no supplemental premium. We show that, within a simple model of consumer demand and insurer plan design, this difference in premium salience can quantitatively explain the bunching in the distribution of MA premiums.

The method we use to tell this story quantitatively is relatively simple, though the exposition is complicated by intricate institutional details. First, we estimate demand and find that consumers are more responsive to the salient than the non-salient premium. Next, we write a simple, highly stylized model of insurer plan design. In the model, insurers facing the actual demand curve, in which consumers are more elastic to the salient premium, respond by setting premiums consistent with the observed distribution. In contrast, when they face a counterfactual demand curve in which consumers respond equally to the two premiums (as if the two premiums were combined), many insurers who previously priced at the bunching point now set substantially lower premiums. Finally, we consider the welfare effect of the information friction. When we "correct" demand so that consumers value the

Figure 1: Actual 2011 Premium Distribution

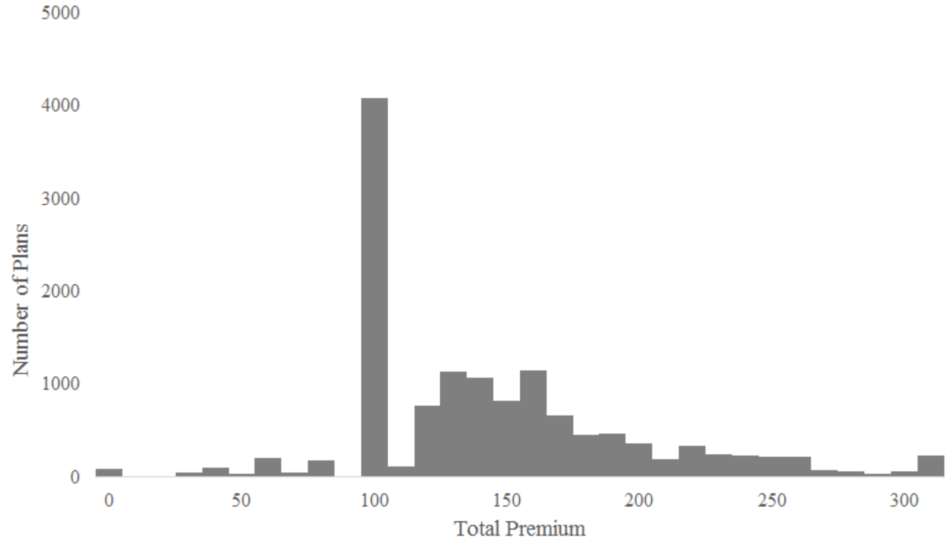

Note: The distribution of total premiums in 2011. The bunching at $96.40 is precisely where insurers are charging the maximum value for the non salient Part B premium and $0 for the salient Part C supplemental premium.

two premiums equally and hold the set of plans constant, consumer surplus increases by about $5/year. In contrast, when we correct demand and also allow insurers to redesign their plans, consumer surplus increases by about $73/year. While we cannot fully rule out alternative explanations for the bunching in the premium distribution, this exercise illustrates the importance of considering the supply-side response when studying behavioral features of consumer demand in health exchanges.

We begin by providing background on Medicare Advantage regulations in Section 2. During our sample period, the nonsalient Part B premium (which pays for the basic package of Medicare coverage) was automatically deducted from Social Security payments and was not displayed on the plan comparison page of the Medicare Plan Finder website.[1] In contrast, the salient Part C supplemental premium (which pays for additional benefits) is paid via check and is advertised as the plan price in most settings, including the Medicare Plan Finder website. In Section 3, we describe our data, which is constructed from two administrative files, one at the plan level and another at the beneficiary level.

In Section 4, we estimate demand, with a particular interest in differences in the consumer

[1] The Medicare Plan Finder website now shows if a plan reduces the Part B premium. We do not currently have data on premiums after this change was implemented, but hope to study the effect in the future.

response to the two premiums. We exploit the rigid premium setting rules in Medicare Advantage to identify a simple logit demand model. Two features provide identification. First, regulations require plans to set initial premiums to reflect the actuarial value of the benefits provided by the plan. Second, based on expected costs relative to traditional Medicare, insurers allocate a rebate which must be used to either buy down premiums or reduce cost sharing requirements. We use before-rebate premiums as a summary statistic for the generosity of a plan, control for insurer and provider network, and use variation in rebate allocation to estimate demand elasticities to plan features. We find that the demand response to the salient supplemental premium is much larger than to the nonsalient Part B premium. In fact, the coefficient on Part B premium reductions in our preferred demand specification is not statistically distinguishable from zero, suggesting that the hidden presentation of the Part B premium affects consumers' ability to optimize. We later corroborate these direct estimates of premium elasticities by showing that our model of insurer plan design can replicate the observed distribution of premiums only given demand primitives similar to the direct estimates.

In Section 5, we perform a modeling exercise to provide an example of the potential importance of the supply-side response for calculating the total welfare effect of information frictions. In the model, profit-maximizing firms face a demand curve and decide how much of their rebate goes to buying down premiums and how much goes to increasing benefits. We run the model with both the estimated demand curve and a series of counterfactual demand curves with varying differences between the two premium elasticities. Imposing equal responses to the salient and nonsalient premiums (as though the two premiums were combined) to predict consumer choices, holding the supply response fixed, we find an upper bound on aggregate gains in surplus of about \$.4/month or \$5/year per beneficiary. However, allowing insurers to respond to these adjust beneficiary preferences generates a much larger gain in consumer surplus—about \$6/month or \$73/year. Here, consumer welfare increases because insurers now lower their nonsalient premiums, effectively filling in the hole to the left of the spike in Figure 1. The exercise illustrates that incorporating a supply response in the analysis of information frictions or choice inconsistencies can dramatically alter calculations of welfare impacts (in this example, by more than an order of magnitude). Though our model makes several assumptions to simplify the exercise, we argue in the text that relaxing these would most likely increase the relative importance of the supply response.

A large body of literature suggests that individuals often make poor choices in the exchange setting. Kling et al. (2012) show that sending individuals a letter containing easily available information about Medicare Part D plans significantly increased the proportion of beneficiaries who switched plans. Polyakova (2016) shows that inertia plays a strong role in Part D plan choice and could prevent consumers from switching plans when a better option arises. Abaluck and Gruber (2011) find the average Part D enrollee could have reduced their expected costs by 23% without increasing the variance in their total costs. Lastly, Afendulis et al. (2014), Bhargava et al. (forthcoming), and Handel (2013) provide "smoking gun" evidence of consumer choice inconsistencies by documenting the choice of strictly dominated plans. This paper provides another example of plan choices that are difficult to rationalize with conventional models and adds the observation that allowing insurers to respond to a hypothetical intervention can induce a much larger welfare change than adjusting consumer behavior alone. This suggests the effects of intervening to correct information frictions could be substantially larger than suggested by estimates that hold the supply-side fixed.

A growing literature studies the supply response to the behavioral frictions. A number of studies examine the supply-side implications of inertia in Part D. A series of papers all broadly find support for an 'invest, then harvest' strategy in which insurers offer low premiums for plans in their first year on the market before gradually raising premiums to exploit the low rates of switching for beneficiaries already enrolled.[2] Starc and Town (2015) note there is limited documentation and appreciation of the idea that competing firms can "de-bias" consumers only when the institutional and regulatory settings make doing so profitable. This paper provides an example of how a friction in premium presentation can distort premium setting, and simulates potential welfare gains from eliminating the friction.

Most closely related is Stockley et al. (2014), who also study the two-premium system in Medicare Advantage. They employ a reduced-form strategy to demonstrate that exogenous shifts in insurer reimbursement rates are passed through to premium reductions only if a plan's total premium is above the Part B premium level. This paper complements their

[2]See Ericson (2014), Ho et al. (2015), Miller (2016), Nosal (2012), and Wu (2015)

work by using a structural approach to separately consider demand and supply. This allows us to simulate counterfactual premiums under hypothetical policy changes and estimate the welfare effect of the information friction, leading to the insight that welfare effects are almost entirely driven by the supply-side response. In addition, we corroborate their findings by reaching similar conclusions on the effect of premium transparency using different methodology and variation.

2 Medicare Advantage Background

Medicare is a major institution in the US health care sector, insuring roughly 50 million people and serving as primary payer for roughly 50% of inpatient medical costs.[3] Traditionally, senior citizens in the United States receive their non-prescription drug health insurance through Medicare Parts A and B. Since 2006, traditional Medicare beneficiaries also have had the option to purchase a separate prescription drug plan provided by a private insurer through Medicare Part D.

In lieu of Parts A and B, Medicare-eligible individuals also have the option of enrolling in a private Medicare Advantage (MA) plan. These plans offer at least as generous coverage (in an actuarial sense; actual plan terms may differ) as traditional Medicare, but have some flexibility to offer coverage beyond that provided by traditional Medicare. About 15 million people were enrolled in Medicare Advantage in 2014, following rapid growth since 2006 when the program had about 7 million enrollees.

The design of MA plans is complex and heavily regulated. Broadly speaking, insurers face three decisions: setting the capitation rate for beneficiaries on a plan; determining supplemental coverages to include in a plan; and allocating a government rebate to improve a plan's cost-sharing or reduce its premiums. While these three decisions are made simultaneously, we present them sequentially for ease of exposition.

Insurers first determine their capitation rate for a plan in a given county. For each county, the Center for Medicare and Medicaid Services (CMS) sets a benchmark reimbursement rate

[3] All facts in this section are drawn from the 2014 Kaiser Family Foundation Medicare and Medicare Advantage fact sheets.

based on how much it costs them to insure the average traditional Medicare patient in that county.[4] Plans then submit an estimate of their cost of providing coverage equivalent to traditional Medicare to an average patient in that county.

The vast majority of plans submit estimates below the benchmark. These plans receive their estimate as capitation and three-quarters of the difference between their estimate and the county benchmark as a rebate. (How the rebate is used is described below.) Plans which submit estimates above the benchmark receive the benchmark as capitation and beneficiaries pay the remainder of the estimate as part of their supplemental premium. The capitation payments ultimately made to insurers are risk-adjusted based on each beneficiary's demographics and medical history.

Insurers then decide what coverage to include in their plan. Coverage can be broken into two distinct segments: basic and supplemental. Plans must offer basic coverage, which includes all of the elements in Parts A and B. The base or Part B premium[5] for basic coverage is set by the Center for Medicare and Medicaid Services (CMS) to be equal to $96.40, which was the Part B premium for traditional Medicare beneficiaries during our sample period.

Supplemental coverage, in contrast, is an optional plan offering and includes elements that are not in Parts A and B. This supplemental coverage may include a Part D plan, dental or vision coverage, home meal delivery in times of sickness, or a number of other options, including improved cost-sharing for areas covered by Parts A and B. Plans set a separate supplemental premium for their supplemental packages. This premium is the sum of the premium for the plan's Part D component (if it has one), which is set by the same process as stand-alone Part D plans, and the premium the plan sets for the other, non-drug, supplemental packages in the plan. Throughout the paper, we will refer to this non-drug supplemental coverage as OS (other supplemental) coverage.

Critical to our identification strategy is that, though it can later be reduced in the rebate allocation process, the before-rebate supplemental premium must be greater than the actuarial cost of the supplemental coverage. This implies that the before-rebate supplemental premium is bounded below by the actuarial cost of the supplemental coverage. The premium

[4]In some counties, CMS sets the benchmark rate above the average cost in order to encourage insurer entry.

[5]This could be altered in the next step, the rebate allocation process.

is also subject to minimum medical loss ratio requirements which ensure that a high percentage of premiums is paid out in benefits. This implies that the before-rebate supplemental premium is bounded above by about (depending on the year) 1.2 times the actuarial cost of the supplemental coverage. Between these two requirements, the before-rebate supplemental premium serves as a fairly precise measure of a plan's supplemental coverage generosity. Since basic coverage is standardized across plans, this also serves as measure of a plan's overall coverage generosity.

Finally, the insurer allocates the rebate that it received in the bidding process. It must use all of the rebate to improve benefits for beneficiaries. The rebate can subsidize the basic premium, the supplemental premium, or cost-sharing in the standard package of benefits. Rebates to cost-sharing are priced actuarially, e.g. a \$500 coinsurance reduction to a procedure that 10% of patients have is considered a \$50 rebate to cost-sharing. The entire rebate allocation process is done behind the scenes by insurers altering the terms of their plan. Consumers do not observe plan rebates, only the final plan terms.

Figure 2 shows the average rebate allocation over time. Insurers allocate most of their rebates to cost-sharing in every year and put more rebate dollars towards the Part D and OS premiums than the Part B premium. Appendix Table A.1 summarizes the plan design information above in a more compact form and provides an example plan.

We can then write the final equations for the basic and supplemental premiums:

$$Basic\ premium = Medicare\ Pt.\ B\ premium - basic\ premium\ rebate$$

$$Supplemental\ premium^{6} = Pt.\ D\ premium + other\ supp.\ premium - supp.\ premium\ rebate$$

These two premiums have different salience because of the way they are paid and presented. The basic premium is not prominently displayed on the Medicare Plan Finder site, as shown in the screenshot in Appendix Figure A.1. In the figure, which shows the initial results from a search query, the highlighted Monthly Premium box shows a \$29 premium for the first plan and a \$0 premium for the second. These numbers represent only the sup-

[6] This formula is slightly altered for the small number of plans that do not provide traditional coverage more cheaply than TM. Beneficiaries on these plans must pay the difference between TM's cost of coverage and the plan's cost of coverage. This amount is added to the supplemental premium.

Figure 2: Average Rebate Allocation Over Time

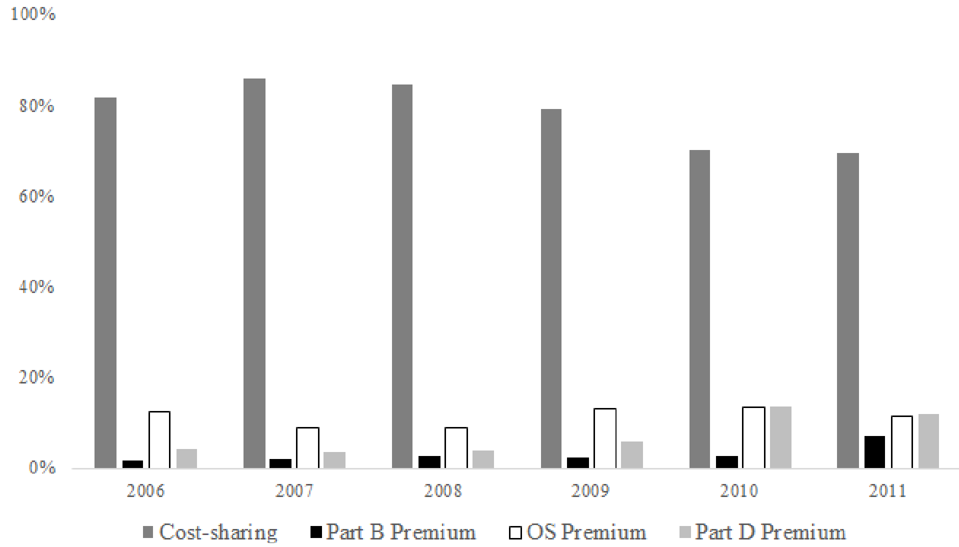

Note: Average rebate allocations (in percent of total rebate) over the sample period. An observation is a plan-county-year. Cost-sharing represents the actuarial value of reductions in plan coinsurance and deductibles. The OS premium covers supplemental packages besides Part D coverage, and the total supplemental premium is equal to the sum of the OS and Part D premiums.

plemental premium; enrollees would also have to pay the Part B base premium for either of these plans. In order to learn the value of the Part B premium, a consumer would have to click on a particular plan and look through a more detailed explanation of benefits. In addition, the Part B premium is generally deducted from Social Security checks, while the supplemental is generally paid out of pocket. Due to both these institutional details, the salience of the two premiums sharply differ. The same differential treatment does not apply to insurers making their rebate decisions; as shown in Appendix Figure A.2, the worksheet insurers fill out to allocate their rebate. The box to enter a Part B rebate is presented symmetrically to cost-sharing and supplemental coverage rebates.

A final detail that will be important to our empirical analysis is that Medicare Advantage insurers must establish a contract with a network of providers before they can offer plans. Based on this contract, they can offer a variety of plans. For example, United Healthcare may enter a contract with Stanford Hospital and Palo Alto Medical Foundation which establishes a relationship and set of reimbursement rates. They could then offer an HMO Basic and HMO Gold plan, both of which feature the same two hospitals as their provider network,

Table 1: Summary Statistics on 2006-2011 Medicare Advantage Plans

Variable	Mean	Standard Deviation	Min	Max	Median
Part D bundled	0.70	0.46	0.00	1.00	1.00
Total rebate	52.85	35.31	0	406.71	49.01
Cost-sharing rebate	40.45	25.80	0.00	222.34	39.47
Supp. premium rebate	10.79	14.04	0.00	245.09	6.60
Supp. premium	41.04	46.76	0	427.2	27.1
Plan enrollment	296.95	1371.74	11.00	72241.00	45.00
Average beneficiary risk score	0.85	0.19	0.35	2.50	0.83
I(rebate Part B > 0)	.034	.17	0	1	0
N	90516				

Note: An observation is a plan-county-year. Part D bundled is an indicator for whether a plan includes Part D coverage. All rebates are in dollars. Plan enrollment is a count of beneficiaries. The beneficiary risk score is computed by CMS and normalized to 1. A risk score of 2 indicates that a beneficiary's expected cost is twice the cost of the average beneficiary.

but which feature different cost-sharing and benefit packages. In this paper, we will use "contract" to refer to a relationship between an insurer and a network of providers and "plan" to refer to a specific insurance plan in which an individual could enroll. Contract IDs enable us to control for insurer-X-provider network-X-county fixed effects, which, to simplify exposition, we will refer to as contract-county fixed effects.

3 Data

We use a panel data set containing enrollment and plan information on every non-employer, non-special-needs, Medicare Advantage plan with over 10 enrollees in the period 2006-2011. Our enrollment figures exclude beneficiaries who are also eligible for Medicaid or who are under 65 but are covered because they meet other, non-age-based criteria. This dataset was constructed using Medicare administrative records. It contains information on risk scores, plan payments, cost estimates, rebate sizes, and rebate allocation decisions, for about 90,000 plan-county-year observations. Those observations are composed of about 12,000 unique plan-years and 5,000 unique MA plans. Table 1 provides a set of summary statistics on the plan-county-year observations.

4 Demand Estimation

4.1 Identification and Specification

Our goal is to separately identify demand coefficients for the base Part B and supplemental premiums. Consider the utility for beneficiary i on plan j as

$$u_{ij} = \delta_{ij} + \epsilon_{ij} \tag{1}$$

where

$$\delta_{ij} = \alpha_1 base\_premium_j + \alpha_2 supp\_premium_j + \phi_j$$

where ϕ_j is plan quality and ϵ_{ik} is i.i.d. and has the standard Type I extreme-value distribution.

In order to estimate α_1 and α_2, we need to address the endogenous correlation between premiums and unobserved elements of plan quality. We employ a rich set of fixed effects and plan characteristics in an attempt to account for quality differences across plans. First, to control for provider network differences, we include contract-county fixed effects,[7] which limits the identifying variation to plans in the same county, offered by the same insurer, with the same provider network.

However, there may still be a correlation between premiums and unobserved benefit quality across plans within the same contract and county. To address this, we control for a rich set of plan characteristics. Our principal control for plan quality is the before-rebate supplemental premium. The Medicare Advantage bidding process requires any coverage beyond the standard benefits covered by traditional Medicare Parts A and B (e.g., lower cost sharing requirements, prescription drug coverage) to be paid for with higher before-rebate premiums. The size of the before-rebate premium therefore directly corresponds to the value of the additional benefits.

Additionally, we are able to split the before-rebate supplemental premium and separately control for Part D and non-drug other supplemental coverage (OS) before-rebate premiums.

[7] Recall that a contract specifies an insurer's relationship with a set of providers, and that an insurer may offer multiple plans on a given contract.

We label the non-drug supplemental and Part D before-rebate premiums as OS and Part D generosity. While we would ideally control for all aspects of the supplemental packages (e.g., dentist visit copay, cost of a new pair of eyeglasses), before-rebate premiums serve as effective summary measures of an otherwise unwieldy contract space.

We additionally exploit the richness of our data to control for any other plan characteristics we observe that may be correlated with quality. Because the basic package of services must be actuarially equivalent to traditional Medicare, it is unlikely there will be significant differences across plans in this standardized set of benefits. There may remain differences in other attributes (e.g., administrative quality), so we control for a plan's CMS star rating, the vintage of the plan, and whether a plan bid above its county benchmark, and if so, how far above the benchmark. Part D prescription drug plans are among the most common and most expensive supplemental packages, therefore we control for whether a plan includes Part D coverage. We also include an indicator for whether a plan includes any non-drug supplemental packages.

While the size of the before-rebate premiums is determined by the plan's estimated value, the final premium is the result of how an insurer allocates a plan's rebate between the two premiums and cost-sharing benefits. We use rebate allocation, conditional on before-rebate premiums, to generate variation in prices. Variation in prices is then driven by, for example, two plans which both have \$50 before-rebate supplemental premiums, but which reduce the \$50 by different amounts in the rebate process.

This approach is necessary in order to use before-rebate premiums to control for quality. Rebates reduce premiums one for one, so measuring the demand response to a change in rebates is the same as measuring the demand response to a change in premiums. This brings us to our estimating equation for the plan attractiveness parameter δ for plan j on contract k in county c at year t:

$$\delta_{ijkct} = \alpha_1 rebate_{jkt}^{base} + \alpha_2 rebate_{jkt}^{supp} + \alpha_3 rebate_{jkt}^{cost}$$

$$+\beta_1 OS\_generosity_{jkt} + \beta_2 PartD\_generosity_{jkt} + \beta_3 I(PartD)_{jkct} + \psi_{kc} + \epsilon_{ijt}$$

where each rebate variable is in dollars and the vector x_{kc} contains further observable plan characteristics, including CMS plan star rating, bundled Part D coverage, and squared terms for Part D and OS generosity and rebates to supplemental premiums and cost-sharing. We do not include a squared term in the Part B premium rebate because limited variation in Part B premiums in the data is likely to prevent us from accurately identifying both a linear and a quadratic effect. Lastly, ψ_{kc} is the contract-county fixed effect. We also include year fixed effects to account for changes in the program and the economy over the sample period.

We run simple logit demand models with traditional Medicare as the outside good and a county-year as a market. Since we ultimately include these demand estimates in a model of supply, we measure plan enrollment as the number of risk-score-weighted months of enrollment, thus accounting for two factors that matter to firms—the risk score weighting of capitation payments and the high mortality of Medicare beneficiaries. To account for heterogeneity in preferences, we use health status to segment beneficiaries into low risk (risk scores below median), medium risk (risk scores between median and 75th percentile), and high risk (risk scores above 75th percentile) samples and estimate demand separately for each segment. A plan's market share in a particular beneficiary segment in a given county is then the percent (including traditional Medicare beneficiaries) of the total risk-months in the beneficiary segment in that county that enroll in the plan.

The set-up is similar to the demand estimation framework in Curto et al. (2015); however, while they include all rebates in a single variable and focus on bids, we are interested in plan design *conditional* on the bid and will consider different types of rebates separately. In addition, we take a different approach to preference heterogeneity by segmenting the beneficiaries.

Our identifying assumption is that rebate allocations are conditionally unconfounded. That is, conditional on before-rebate premiums and other plan characteristic controls, differences in rebate allocation between plans offered on the same contract in the same county are uncorrelated with unobservable elements of the plans.

There are two main threats to this identification strategy. First, the rebate allocation decision may be correlated with what plan packages are offered. For example, an insurer may target its basic plan to healthy beneficiaries and its premium plan to sick beneficiaries. It would then possibly apply the rebate to premiums for the basic plan and include a gym membership as a supplemental package; whereas, for the premium plan, it may apply the rebate to cost-sharing and include assistive bars in bathrooms as a supplemental package. Our strategy will be more effective if financial features of plans are more important for consumer decisions than other factors.

Second, plans within a contract may have different levels of total rebate to allocate in the first place. This could occur because of variation over time in bidding or benchmarks, or if projections for utilization across the two plans are different and lead actuaries to compute a different cost estimate. If this is because of exogenous shifters, such as changes in the CMS benchmarks or in the bargaining landscape with providers, it will not affect the validity of our estimates. If it is because of a contract drastically changing its provider network to cut costs, we will conflate a positive change in rebate allocation with a restrictive change in provider network. Unfortunately, we do not observe changes to provider networks within a contract.

4.2 Variation

The variation in plan characteristics necessary to identify our model generally arises from the quality tiers of an insurer's product offering. For example, Blue Cross offered both an HMO Standard and an HMO Plus plan in Boise, Idaho from the same contract from 2006-2011. Both of these plans featured the same provider network, but they had different supplemental packages and financial features. The standard plan rebated Part B premiums for the middle two years of the sample, but not the other four years, while the Plus plan added a Part D premium rebate in its second year.

This example is not necessarily representative, and Appendix Table A.2 contains more systematic statistics on within-contract variation. The median contract has a range of $23 between its most and least generous supplemental packages, as well as a range of $18 between its largest and smallest rebate to cost-sharing.

Given the infrequency of plans offering rebates to Part B premiums, the contracts which offer at least one such plan are of particular importance for identifying α_1. We are relying on within-contract variation to isolate demand responses to plan features other than provider networks so it is critical to have sufficient variation within the contracts which offer plans

Table 2: Characteristics of Plans that Rebate Part B Premium

Mean of:	Plans w/Part B rebate > 0	Plans w/Part B rebate= 0
Total supp. premium	22.15	52.81
I(Part D)	.63	.71
I(HMO)	.56	.30
Total rebate	118.12	50.75
Plans in county	34.8	28.2
MA share in county	.21	.18
Plans in contract	2.83	2.17
N	3078	87438

Note: An observation is a plan-county-year. Column 1 shows the mean of a variety of plan characteristics for plans that rebate some of the Part B premium, while Column 2 shows the same statistic for plans that rebate none of the Part B premium. I(Part D) is an indicator for if Part D coverage is bundled with a plan. I(HMO) is an indicator for if a plan is HMO. Plans in county is the total number of plans active in a county-year. MA share in county is the proportion of Medicare beneficiaries in a given county-year enrolled in Medicare Advantage. Plans in contract is the total number of plans offered on the same contract as a given plan. Total supplemental premium and total rebate are in dollars.

rebating Part B premiums. Some statistics on these plans are given in Table 2. These contracts, and plans, are generally in more crowded Medicare Advantage markets. Plans offering Part B rebates are in counties with 34.8 plans on average, while plans not offering Part B rebates are in counties with 28.2 plans on average. They are less likely to offer prescription drug coverage (63% versus 71% for plans without Part B rebates), more likely to be an HMO (56% versus 30%), and more likely to be offered within a provider contract which includes more plans (2.83 versus 2.17). We control explicitly for the former two differences, and the additional plans within the provider contract allow us to control for the effects of the provider network. Plans offering Part B rebates also tend to offer less generous supplemental coverage, with just 22.15 worth of coverage on average compared to 52.81 for the average plan without Part B rebates. There just over 3,000 plan-years offering Part B rebates, representing 3.4% of total plan-years.

4.3 Results

Results from the primary specification are shown in Table 3. Columns 1-3 show the results for the healthy, medium, and sick segments of the population, respectively, and Column 4 shows the results for the aggregate sample. Focusing first on the aggregate sample, rebates to the supplemental premium are strongly significant, with a magnitude implying that if a plan with a 1% market share (around the 80th percentile of plan market share) were to increase its supplemental premium rebate from $0 to $50, its market share would increase to about 1.5%. This is a relatively large increase, although it is small in absolute terms because traditional Medicare holds the lion's share of most markets. Rebates to Part B premiums, on the other hand, have an effect that isn't significantly different than 0, while rebates to cost-sharing have a significant effect that is slightly larger than the supplemental premium effect.

In the aggregate sample, elasticity to OS generosity is positive, while the elasticity to the Part D generosity is negative. Recall that our identification strategy only requires that the *rebate allocation* is uncorrelated with plan quality, and actually uses the generosity measures as controls for quality. Thus, in our model, these coefficients are a combination of price and quality effects, implying that standard demand model concerns about positive "price" elasticities do not map to our context.

Turning now to the results for the different segments of beneficiaries, we see that the lowest risk individuals (Column 1) have the largest coefficient on supplemental premiums and the lowest coefficient on cost-sharing rebates. In contrast, in the highest risk segment, the coefficient on the supplemental premium rebate is not significant, but the sensitivity to cost-sharing rebates is higher than in the other segments. It also seems that risky beneficiaries have a relative preference for plans with higher supplemental premiums (indicating more extensive coverage). As in the overall sample, none of the three segments have a statistically significant responses to the nonsalient Part B premium rebate.

As a robustness check, note that the supplemental premium is the sum of the Part D premium and the OS premium. Rebates to the supplemental premium can be split into rebates to these two separate premiums. In our main specification, we include only the sum

Table 3: Base Demand Specification

D.V. = log market share ratio	Low-Risk β (S.E.)	Medium-Risk β (S.E.)	High-Risk β (S.E.)	Aggregate β (S.E.)
Part B rebate	.32	.052	.36	.11
	(.37)	(.42)	(.47)	(.41)
Supp. premium rebate	1.28***	1.01**	.76	.98**
	(.37)	(.45)	(.46)	(.42)
Supp. pmm. rebate^2	.0014	.0034	.0046	.0034
	(.005)	(.0057)	(.006)	(.0055)
Cost-sharing rebate	1.13***	1.21***	1.35***	1.22***
	(.28)	(.3)	(.33)	(.3)
Cost-sharing rebate^2	-.005**	-.0043*	.0038	-.004*
	(.0022)	(.0024)	(.0026)	(.0024)
Part D bundled	98.4***	86.1***	78.3***	85.2***
	(23.9)	(25.5)	(26.7)	(25)
I(Supp. coverage)	-34.1***	-36.2***	-37.7***	-35.7***
	(7.25)	(8.07)	(9.48)	(8.12)
Total Part D premium	-1.34**	.96	.67	.97
	(.66)	(.7)	(.71)	(.67)
Total Part D premium^2	.0047	.0036	.0045	.0049
	(.0059)	(.006)	(.006)	(.0058)
Total OS premium	.015	.57**	1.09***	.65**
	(.26)	(.28)	(.29)	(.27)
Total OS premium^2	.00074	.00045	-.00017	.00038
	(.0014)	(.0016)	(.0016)	(.0015)
Constant	-446.4***	-468.6***	-496.3***	-471.5***
	(17.5)	(18.7)	(18)	(17.4)
N	90,061	89,187	88,066	90,070

$^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$

Note: Low-risk is individuals with a CMS risk score below the median. Medium-risk is individuals with a risk score in the third quartile. High-risk is individuals with a risk score in the highest quartile. Demand is estimated separately for each of these segments with traditional Medicare as the outside option. All rebate variables and premiums are taken in dollars. The supplemental premium rebate is the sum of the rebates directed to the Part D premium and the OS premium. Part D bundled is an indicator for whether Part D coverage was bundled with the plan, and I(Supp. coverage) is defined in the same way. Total Part D premium is the sum of the premium paid and rebate given for Part D coverage, and total OS premium is defined in the same way. The sample spans the time period 2006-2011. Observations change across risk segments because we censor plans with less than 10 beneficiaries in each segment.

of these rebates, but we separate the categories in Appendix Table A.3 . The Part D rebate has a larger coefficient than does the combined supplemental rebate in our main specification, while the coefficients on the OS rebate are substantially smaller and feature large standard errors, suggesting the Part D rebates are important drivers for our demand estimation.

In summary, the results indicate that choice responds to reductions in the supplemental premium, but not to reductions in the Part B premium. While we interpret this pattern as a product of the differential presentation of the two premiums, we cannot rule out that consumers have truly different preferences over premium dollars paid for supplemental coverage as compared to dollars paid for Part B coverage. This could arise, for instance, if the consumer response to reductions in premiums above $96.40 is substantially different from the response to reductions below $96.40, and if this response changes in a way that is not captured by our linear and quadratic terms.

5 Insurer response to information friction

5.1 Model

Differences in the salience of different premiums can distort not only the plans individuals choose, but also the plans that insurers choose to offer. In this section, we specify a simple model of plan design that shows that differential responses to salient and nonsalient premiums, consistent with our demand estimates, can explain the bunching in the premium distribution in Figure 1. The goal of this exercise is to provide a plausible example of the relative importance of the supply-side response when addressing information frictions and choice inconsistencies in contexts like marketplaces for health insurance.

In the model, insurers choose a rebate allocation in order to maximize profit, holding all other plan characteristics fixed. An attractive feature of the model is that is allows us to simulate rebate allocations using only parameters from the demand system, with no estimation of additional supply-side parameters. Insurers face the logit demand system estimated in the previous section.

To keep the example as simple and transparent as possible, we only allow insurers to adjust their rebate allocation, holding plan attributes (such as bids, supplemental packages, and before-rebate premiums for supplemental packages) constant. This reduces the firm problem to one parameter. We also assume that multi-plan firms will not jointly consider the rebate allocations of separate plans.[8] These restrictions likely reduce the size of the supply response by restricting the number of margins that insurers can adjust, as we discuss in detail when reporting the model results.

For a given rebate allocation, plan j in market segment k has logit parameter:

$$\delta_j^k = \alpha_1^k * rebate\_PartB_j + \alpha_2^k * rebate\_supp_j + \alpha_3^k * rebate\_PartB_j^2 + \alpha_4^k * rebate\_supp_j^2$$

$$+\alpha_5^k * rebate\_cstshr_j + \alpha_6^k * rebate\_cstshr_j^2 + \xi^k * X_j + \epsilon_{jk}$$

where all plan characteristics besides rebate allocations are in the vector X_j. Each plan's total (across market segments) risk-weighted enrollment is then given by:

$$E_j = \sum_{k \in 1:3} M_k \frac{\delta_j^k}{\sum_i \delta_i^k} \tag{2}$$

where M_k is the total amount of risk-weighted enrollment available in market segment k. Under perfect risk adjustment, the difference between capitation payments and expected costs from each risk-month is identical. Therefore, actuarial requirements on rebates to cost-sharing imply that rebate allocation decisions do not affect the profitability from a given beneficiary. Under these assumptions, maximizing profit is equivalent to maximizing total risk-months of enrollment E_j, allowing us to simulate rebate allocation decisions without estimating costs.

A Nash equilibrium in this model is a vector of rebate allocations X which maximizes Equation 2 for every plan. We allow plans to choose the share of their rebate to allocate to premiums on a grid with step size 1%.[9] We use data from 2011 and iterate over firms' best response functions from the starting point of all firms devoting all of their rebate to

[8] This assumption greatly reduces computational burden and makes little difference in our estimation, as over 99% of plans have a strictly dominant strategy.

[9] With the assumption that plans rebate all of the supplemental premium before rebating any of the Part B premium (which is true in the data), determining what percent of rebate goes to premiums is sufficient to fully determine the rebate allocation.

cost-sharing until we arrive at a mutual best response.

The intuition for the insurer's decision is quite simple. The insurer should allocate the marginal rebate dollar to the category which has the largest marginal effect on total demand. The squared terms in the demand system allows for increasing or decreasing returns to allocating rebate dollars to a given category, as otherwise each plan would allocate their entire rebate to one category. The marginal effect of additional rebates to category r can be written as:

$$\frac{dE}{d(rebate\_r)} = M_1 \frac{dE^1}{d(rebate\_r)} + M_2 \frac{dE^2}{d(rebate\_r)} + M_3 \frac{dE^3}{d(rebate\_r)}$$

The marginal return to category r is a function of two things: the estimated elasticity to the category and the risk score composition of a given county. A higher estimated elasticity for category r will increase the likelihood that it is the most effective place for the marginal dollar. This is the channel through which salience effects filter to the supply-side—lower demand elasticity to the Part B premium makes insurers less likely to devote dollars to that category.

Different demographic mixes across markets can also generate substantial differences in firm behavior. To give an extreme example, if a plan were in a market with only sick beneficiaries, they would devote the first \$45 of their rebate to cost-reductions before allocating any rebate to premium reductions. On the other hand, plans in a market with only the healthiest beneficiaries would devote all of their rebates to premium reductions for any sensible value of rebate.

To explore whether our estimated demand parameters best explain the distribution of premiums under this model of supply, we also consider several alternative demand systems. We consider $\alpha_1^k = \gamma\alpha_2^k, \alpha_3^k = \gamma\alpha_4^k$ for a variety of $\gamma \in [0, 1]$, where γ provides a measure of the relative demand response to the different premiums. When $\gamma = 1$, consumers respond equally to the two premiums. When $\gamma = 0$, consumers do not respond at all to the less salient premium. This approach allows us to rely less on the precision of our demand estimates, as we can consider the supply-side implications for a variety of demand elasticities within our confidence interval. We consider each γ on a grid from 0 to 1 with step size .1. Our demand

Table 4: Results from Supply Model

γ	% at bunch point	% below bunch point	Model MSE
0	48.8	0	586.9
0.1	48.2	1.1	569.3
0.2	47.2	2.1	544.3
0.3	46.3	3.0	529.2
0.4	45.0	4.3	546.3
0.5	42.0	6.4	603.5
0.6	39.2	10.1	703.1
0.7	33.0	16.3	887.4
0.8	21.8	27.5	1198.8
0.9	9.6	39.7	1499.0
1.0	0	49.2	1601.6
Actual	29.7	5.7	NA

Note: An observation is a plan-county. γ is the salience of the Part B premium relative to the supplemental premium. % at bunch point is the percent of plan-counties with total premium 96.4 in the data and the percent with total premium - 96.4 $<$ .01*(total rebate) in the simulation results. % below bunch point is the percent of plan-counties with total premiums below 96.4 in the data and the percent below 96.4-.01*(total rebate) in the simulations. MSE is the mean squared error. Year is 2011.

estimates correspond to $\gamma = 0$.

5.2 Results

The results of this estimation, shown in Table 4, show that the difference in premium salience is important enough to generate the observed bunching in the premium distribution. With $\gamma = 0$, approximately the estimated value, we predict 48% of plans at the bunch point and no plans below it. The actual premium distribution has 29.5% of plans at the bunch point and 5.7% below it. Setting $\gamma = .5$ better predicts the percent at the bunch point (42), but generates too many plans below the bunch point (6.4%). $\gamma = .3$ achieves the lowest mean squared error.

While not a conclusive explanation for the bunching, this exercise shows that information frictions like the MA two-premium system could be important enough to generate meaningful supply-side effects. We cannot rule out alternative explanations such as an anchoring effect at the Part B premium, but we show that these additional forces are not necessary to explain the bunching.

As mentioned earlier, these results provide strong support for differential demand elasticities. Observed plan design should be optimal given the underlying demand system, and the rebate allocations observed in the data are only optimal if there are different elasticities to the two premiums.

In addition to suggesting the presence of a meaningful supply-side response, we can also predict counterfactual premium distributions. Setting $\gamma = 1$ simulates a policy that makes the two premiums equally salient (as if the Part B and supplemental premiums are presented as a single premium). Figure 3 compares the actual distribution of premiums to the simulated distribution with $\gamma = 1$, and $\gamma = .7$. The distribution with $\gamma = 1$ suggests that if premiums were presented identically, firms would offer a much smoother distribution of premiums. This counterfactual plan menu offers many more low-premium plans than we actually observe in the data, which substantially improves consumer welfare, as we show in the next section.

There are two reasons we may underestimating the magnitude of the insurer shift to lower premium plans. Recall that we do not allow firms to alter their bids, supplemental packages, or before-rebate supplemental premiums. This assumption has important implications for our simulation of a policy that equalizes salience of the two premiums (setting $\gamma = 1$). First, it may be that insurers would actually lower their bids under this policy because rebate dollars are more valuable. Second, insurers may also offer less generous supplemental benefits in this scenario. This could occur if insurers were previously adding supplemental packages and then rebating the additional premium, and now prefer to forgo those packages and reallocate rebate dollars to the Part B premium. Both of these factors would increase total rebates and likely strengthen the shift towards low-premium plans that occurs in our simulation, underscoring that the supply-side response is crucial for the welfare consequences information frictions and choice inconsistencies in consumer plan choice.

5.3 Consumer Welfare

We use the predictions of our supply model to estimate the equilibrium effect of the information friction on consumer welfare. We additionally decompose the share of the change

Figure 3: Comparisons of Simulated and Actual Premium Distributions

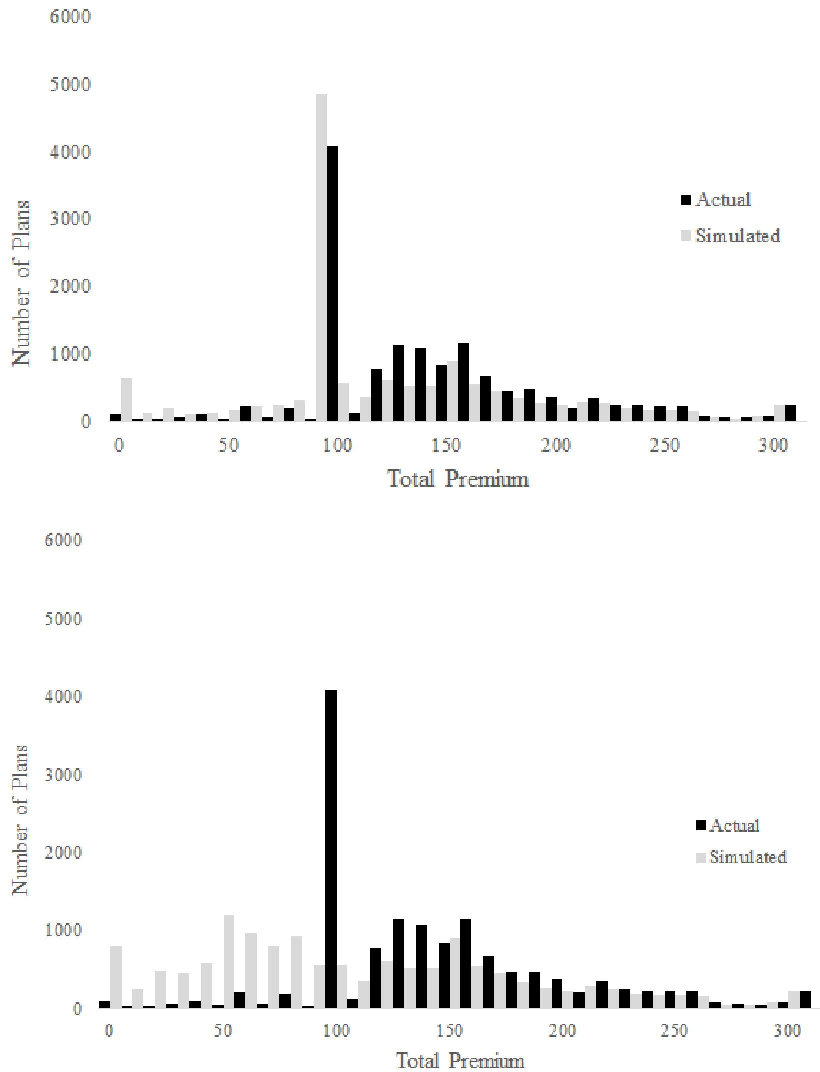

Note: The top panel compares actual premiums and model predicted premiums with $\gamma = .7$ (meaning the demand response to the Part B premium is 70% of the response to the supplemental premium). The bottom panel shows predicted premiums with $\gamma = 1$ (as if the two premiums were combined).

in consumer welfare due to distorted choices and the share due to supply responses.

We estimate the effect on consumer welfare by comparing consumer surplus when consumers make choices with $\gamma = 1$ (consumers consider the two premiums jointly) versus $\gamma = 0$ (consumers consider only the salient supplemental premium). We make this comparison in two steps. First, we fix the set of plans available and compute the change in welfare generated by moving γ from 0 to 1. This isolates the welfare impact of distortions in choices. Second, we fix γ at 1 and compute the change in welfare generated by changing the available plans from the observed set to the set our supply model predicts for γ=1. This is the portion of the change in welfare due to the supply response.

5.3.1 Welfare effects from distortion of choices

To estimate the effect of the salience distortion on consumer choices, we want to hold the set of plans constant and estimate how utility changes when consumers chose plans with $\gamma = 0$

$$u_{ij} = 0 * rebate\_PartB_j + \beta_2 rebate\_supp_j + \beta_3 rebate\_supp_j^2 + \xi X_j + \epsilon_{ij}$$

versus $\gamma = 1$:

$$u_{ij} = \beta_2(rebate\_PartB_j + rebate\_supp_j) + \beta_3(rebate\_PartB_j + rebate\_supp_j)^2 + \xi X_j + \epsilon_{ij}$$

While we can easily run these choice models and predict plan enrollment, the standard logit machinery for determining welfare does not apply here, because when $\gamma = 0$, consumers do not necessarily choose the plan that gives them the highest utility. The standard method does not capture the surplus consumers still receive from Part B premium rebates because they do not consider it when making their plan choice.

We address this problem by putting an upper bound on the welfare gain. We estimate the number of additional people[10] who choose plans that rebate Part B premiums, as well as the change in the total number of dollars rebated, when $\gamma = 1$ instead of 0. The change in total Part B rebate collected provides an upper bound for total surplus gained, as it is the

[10] We convert from risk-months to people by dividing risk-months by the mean number of risk-months per person per year in each of our three market segments.

Table 5: Consumer Surplus Results

Upper bound on welfare gain ($/person/month) of moving from $\gamma = 0$ to $\gamma = 1$ without supply response

	Market segment			
Population	Healthy	Medium	Sick	Aggregate
Switchers	72.19	75.15	77.86	74.45
Aggregate	.35	.38	.44	.38

Welfare gain ($/person/month) of supply response to moving from $\gamma = 0$ to $\gamma = X$

	Market segment			
γ	Healthy	Medium	Sick	Aggregate
0.1	0.006	0.004	0.001	0.004
0.5	0.28	0.26	0.17	0.25
0.7	2.36	2.29	1.53	2.15
0.8	4.99	4.51	2.43	4.27
1	7.94	6.35	1.75	6.10

Note: The differences in surplus are in dollars per person per month. The top panel considers only the effect of changes in consumer choices, holding supply fixed, and represents an upper bound on the increase in surplus. The lower panel includes a supply response, and also drops the 5% of counties with largest maximum premium rebate.

surplus an individual would gain if she were indifferent between her choice with $\gamma = 0$ and her choice with $\gamma = 1$ before taking Part B rebates into account.

We should expect to see an increase in consumer surplus, because agents now take the rebates to Part B premiums into account in their choice. Some agents will realize that plans that rebate some of the Part B premium are actually their preferred choice and switch to a plan that gives them more utility. Since we hold the set of plans constant, no consumers should lose utility since they can always stay enrolled in their current plan.

The results are reported in the first panel of Table 5. The gains among those who switch plans are large: around $70 per person per month on average. Few consumers switch, however, so the average gains in the population are much smaller, less than a dollar per month.

5.3.2 Welfare effects from distortion of plan design

To estimate the welfare effect from the information friction's effect on supply, we first fix consumer utility to have $\gamma = 1$:

$$\begin{aligned} u_{ij} = & \beta_2(rebate\_PartB_j + rebate\_supplemental_j) \\ & + \beta_3(rebate\_PartB_j + rebate\_supplemental_j)^2 + \xi X_j + \epsilon_{ij} \end{aligned}$$

For the supply-side welfare estimates, we assume that this is true consumer utility and determines plan choice. To compute consumer surplus, we compare utility for plan choices under each set of plans generated by the supply-side model for different assumed values of γ. This returns the share of the equilibrium welfare effect attributable to the supply-side responses to the information frictions (i.e. the welfare impact resulting from changes in plan design when insurers incorporate different values of γ).

For this exercise, we can use the standard logit welfare estimation tools and analyze the effect of changing the set of plans available to consumers. The dollar difference in average consumer surplus between two sets of plans in a market is given by Train (2009):

$$\Delta E(CS) = \frac{1}{\alpha}[log(\sum_{j=1}^{n} exp(\delta_j^1) - log(\sum_{j=1}^{n} exp(\delta_j^0)]$$

where α is the price elasticity in utility and δ_j is the logit parameter for plan j. Using the distribution of plan characteristics generated by the model of supply, we compute logit parameters for each plan under different assumptions about γ and then use the above formula to compute the dollar-valued change in surplus.

Some consumers will see an increase in surplus in this case because the menu of plans available changes to their benefit. In particular, turning again to the premium distributions shown in Figure 3, insurers will offer substantially more low-premium plans. However, this shift could lead to a decrease in surplus for high-risk consumers if they prefer plans with higher benefits to plans with low premiums.

In our preferred results, shown in the bottom panel of Table 5,[11] we find that the aver-

[11] In our preferred specification, we drop the 5% of counties with the largest maximum premium rebate.

age gain in consumer surplus from changing the set of available plans is $6.10 per month. Intuitively, when we break up surplus gains by patient characteristics, the average gain is $7.94 for the healthy population, $6.35 for the medium, and $1.76 for the sick. These gains in surplus are vastly larger than our estimates of the surplus gained by eliminating the choice distortion with the set of plans fixed, highlighting the importance of accounting for the supply-side response when considering the impact of information frictions or consumer choice inconsistencies. Focusing only on the consumer response ignores a substantial portion (in this example, the majority) of the equilibrium effect, and could give a misleadingly low estimate of the total loss of consumer surplus.

Additional heterogeneity across counties is shown in Figure 4, which plots the average surplus gained by each consumer market segment in each county. Despite the average gain in surplus for the sick population nationally, they are clear losers in some counties. In fact, the sickest segment of the population sees an average decrease in surplus in 44.9% of counties. The middle segment sees an average decrease in about 1% of counties, while the healthy segment never sees a decrease. As expected, the healthiest beneficiaries typically reap the largest surplus gains, as the simulations generally increase the number of low premium plans available.

Total Medicare enrollment, net of dual-eligible beneficiaries, in 2011 was 37.6 million,[12] and the lower bound on the increase in consumer surplus by eliminating the difference in salience between premiums is $6.10 per enrollee per month. This implies the total loss of consumer surplus due to not displaying the Part B premium to consumers is as large as $229 million per month. Curto et al. (2015) provide another benchmark. They estimate the total consumer surplus from Medicare Advantage to be about $50 per beneficiary-month, implying that correcting premium presentation could be on the order of 10% of the gains from enrolling in Medicare Advantage for the entire Medicare population (traditional Medicare and MA enrollees).

Before correcting for outliers, we find that increasing γ from 0 to 1 increases consumer surplus by an average of $18.35. However, this figure is driven by an extremely long right tail. To prevent these plans from driving our results, we prefer to use winsorized means or full distributions.

[12]Kaiser Family Foundation Medicare Enrollment report

Figure 4: Average Surplus Change Across Counties

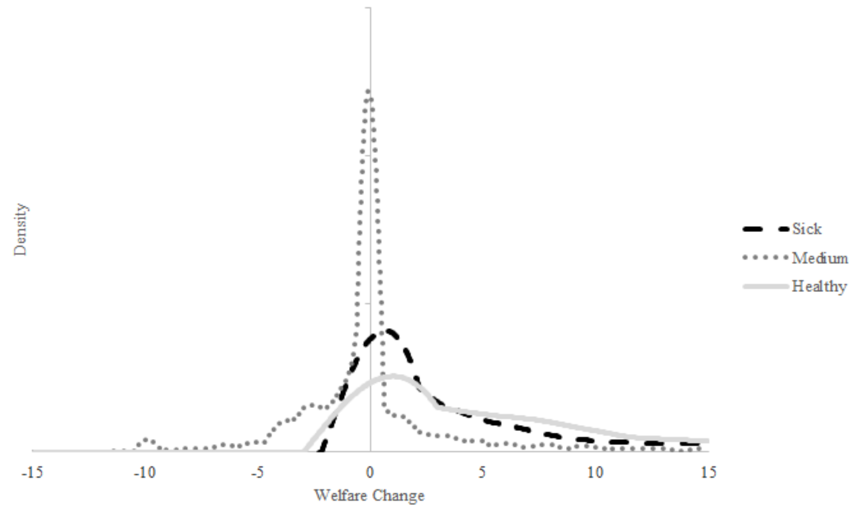

Note: This figure shows the county-level kernel density of the change in monthly average consumer surplus induced by moving γ from 0 to 1. This is equivalent to a reform that combines the Part B and Part C supplemental premiums into a single premium. The different lines represent the surplus change for different beneficiary segments—"healthy" is beneficiaries with below median risk score, "medium" is the third quartile, and "sick" is the top quartile.

6 Conclusion

We examine the effects of a distortion in the presentation of premiums on Medicare Advantage insurance exchanges. With the trajectory of health care reform in the United States increasingly coming into focus as a system built around publicly-operated private insurance exchanges, getting the design of these systems right must be a central focus. A number of articles have documented choice distortions in these markets attributable to information frictions, choice overload, inertia, and unsophisticated consumers. In this article, we document a large price salience effect in the Medicare Advantage insurance marketplace and show that salience differences generate very different demand elasticities for the two premiums MA beneficiaries pay.

The effect of such information frictions or choice inconsistencies on the plans insurers choose to offer has received less attention. We estimate a simple supply model to test whether the demand elasticities we estimate can explain bunching in the distribution of Medicare Advantage premiums. Differences in premium demand elasticities close to our estimates are necessary to explain the bunching and the small number of plan offerings below the bunch point.

We then examine welfare by simulating the consumer choices that would arise if the premiums were equally salient. When we hold the menu of plans available constant and allow consumers to re-optimize their plan choice, we find a modest average increase in surplus of $.4/month or $5/year per Medicare Advantage enrollee. The increased surplus is concentrated among those switching plans, who receive a $600/year bump in surplus when the information friction is eliminated.

The largest effects, however, come when we change the set of available plans according to our model's prediction and then allow consumers to re-optimize. Under this counterfactual, average consumer surplus increases by about $6/month or $73/year, an order of magnitude larger than consumer re-optimization with the choice set held constant. This estimate requires the considerable assumption that changing premium presentation affect insurer decisions only through their rebate allocations, but there are generalizable lessons to be gleaned from this exercise.

This paper provides evidence that information frictions and choice inconsistencies can have large effects on insurer plan offerings, which in turn can have large effects on consumer surplus. As the market for health insurance increasingly moves to the types of exchanges analyzed here, designing these markets becomes a critical public policy problem. Institutional details that affect the choice decision of consumers can have large effects on their plan choices and the plans offered on exchanges. Understanding the interactions between exchange design, consumer demand, and insurer plan offerings is an important area for continued research.

Bibliography

Abaluck, J., and J. Gruber. (2011). “Heterogeneity in choice inconsistencies among the elderly: Evidence from prescription drug plan choice.” *The American Economic Review* 101.3: 377-381.

Afendulis, C. C., Sinaiko, A. D., and Frank, R. G. (2015). “Dominated choices and Medicare Advantage enrollment.” *Journal of Economic Behavior and Organization.* 119: 72-83.

Bhargava, S., Loewenstein, G., and Sydnor, J. (2017). “Evaluating Health Insurance Decisions: Health Plan Choices from a Menu With Dominated Options.” *The Quarterly Journal of Economics, forthcoming.*

Cabral, M., Geruso, M., and Mahoney, N. (2014). “Does privatized health insurance benefit patients or producers? Evidence from Medicare Advantage.” *NBER Working Paper w20470.*

Cawley, J., Chernew, M., and McLaughlin, C. (2005). “HMO participation in Medicare+ Choice.” *Journal of Economics & Management Strategy.* 14.3: 543-574.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2014). “Can Managed Competition Work? The US Medicare Advantage Program.” Mimeo.

Duggan, M., Starc, A., and Vabson, B. (2016). “Who benefits when the government pays more? Pass-through in the Medicare Advantage program.” *Journal of Public Economics.* 141: 50-67.

Ericson, K. M. (2014). “Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange.” *American Economic Journal: Economic Policy.* 6.1: 38-64.

Finkelstein, A. (2009). “E-ztax: Tax salience and tax rates.” *The Quarterly Journal of Economics.* 124.3: 969-1010.

Handel, B. R. (2013). “Adverse selection and inertia in health insurance markets: When nudging hurts.” *The American Economic Review.* 103.7: 2643-2682.

Heiss, F., D. McFadden, and J. Winter. (2010). “Mind the gap! Consumer perceptions and choices of Medicare Part D prescription drug plans.” *Research findings in the economics of aging.* University of Chicago Press: 413-481.

Ho, K., Hogan, J., and Morton, F. S. (2015). “The impact of consumer inattention on insurer pricing in the Medicare Part D program." *NBER working paper w21028.*

Kling, J. R., Mullainathan, S., Shafir, E., Vermeulen, L. C., and M.V. Wrobel. (2012). "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics*, 127.1.

Miller, K. (2016). "Do Private Medicare Firms Have Lower Costs?" Unpublished manuscript.

Newhouse, J. P., and McGuire, T. G. (2014). "How Successful Is Medicare Advantage?" *Milbank Quarterly.* 92.2: 351-394.

Nosal, K. (2012). "Estimating switching costs for Medicare Advantage plans." Unpublished manuscript.

Polyakova, M. (2016). "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics.* 8.3: 165-195.

Song, Z., Landrum, M. B., and Chernew, M. E. (2013). "Competitive bidding in Medicare Advantage: Effect of benchmark changes on plan bids." *Journal of Health Economics.* 32.6: 1301-1312.

Starc, A., and Town, R. J. (2015). "Internalizing Behavioral Externalities: Benefit Integration in Health Insurance." *NBER working paper w21783.*

Stockley, K., McGuire, T., Afendulis, C., Chernew, M. (2014). "Premium Transparency in the Medicare Advantage Market: Implications for Premiums, Benefits, and Efficiency" *NBER working paper w20208*

Wu, Y. (2016). "Supply response to consumer inertia: strategic pricing in Medicare Part D" *Doctoral dissertation, Massachusetts Institute of Technology.*

A Appendix Figures and Tables

Figure A.1: Premium Display on Medicare Plan Finder Website

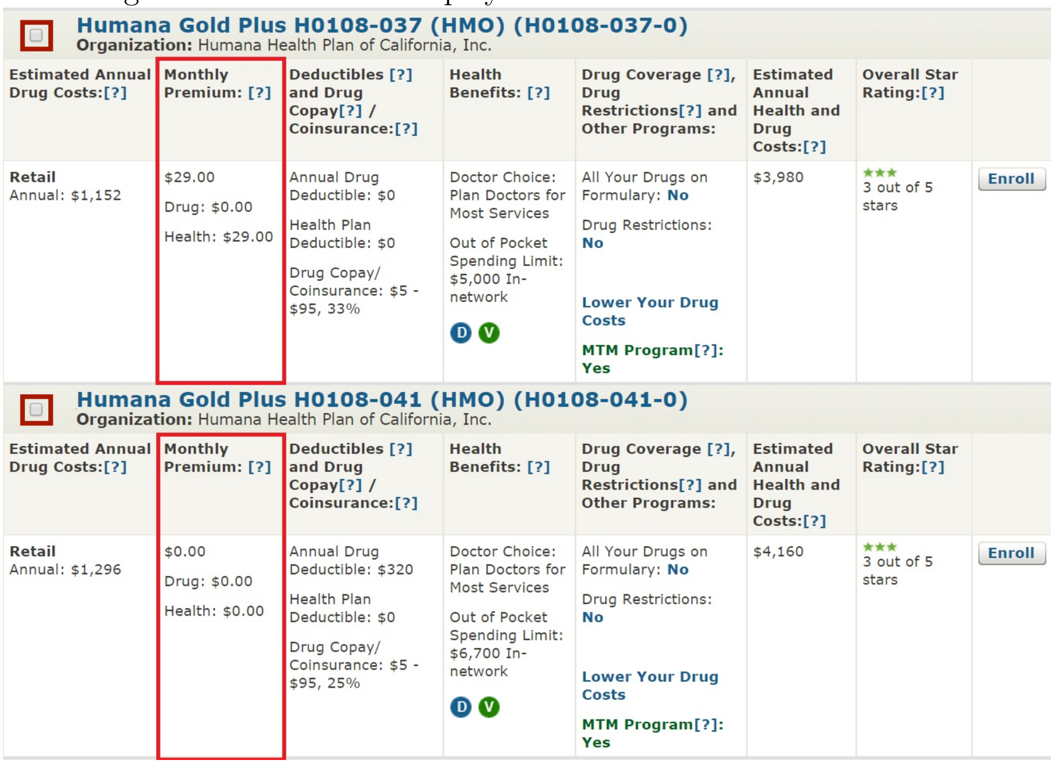

Humana Gold Plus H0108-037 (HMO) (H0108-037-0)
Organization: Humana Health Plan of California, Inc.

Estimated Annual Drug Costs:[?]	Monthly Premium: [?]	Deductibles [?] and Drug Copay[?] / Coinsurance:[?]	Health Benefits: [?]	Drug Coverage [?], Drug Restrictions[?] and Other Programs:	Estimated Annual Health and Drug Costs:[?]	Overall Star Rating:[?]	
Retail Annual: $1,152	$29.00 Drug: $0.00 Health: $29.00	Annual Drug Deductible: $0 Health Plan Deductible: $0 Drug Copay/ Coinsurance: $5 - $95, 33%	Doctor Choice: Plan Doctors for Most Services Out of Pocket Spending Limit: $5,000 In-network D V	All Your Drugs on Formulary: No Drug Restrictions: No Lower Your Drug Costs MTM Program[?]: Yes	$3,980	3 out of 5 stars	Enroll

Humana Gold Plus H0108-041 (HMO) (H0108-041-0)
Organization: Humana Health Plan of California, Inc.

Estimated Annual Drug Costs:[?]	Monthly Premium: [?]	Deductibles [?] and Drug Copay[?] / Coinsurance:[?]	Health Benefits: [?]	Drug Coverage [?], Drug Restrictions[?] and Other Programs:	Estimated Annual Health and Drug Costs:[?]	Overall Star Rating:[?]	
Retail Annual: $1,296	$0.00 Drug: $0.00 Health: $0.00	Annual Drug Deductible: $320 Health Plan Deductible: $0 Drug Copay/ Coinsurance: $5 - $95, 25%	Doctor Choice: Plan Doctors for Most Services Out of Pocket Spending Limit: $6,700 In-network D V	All Your Drugs on Formulary: No Drug Restrictions: No Lower Your Drug Costs MTM Program[?]: Yes	$4,160	3 out of 5 stars	Enroll

Note: A screen shot of the Medicare Plan Finder site in 2014. The monthly premium displayed includes only the supplemental premium. The Part B premium is not displayed on the initial search results page.

Figure A.2: Rebate Allocation Worksheet

B. MA Rebate Allocation

	Rebate PMPM Allocation				Maximum
	Medical	Admin	Gain / (Loss)	Total	Value
1. MA Rebate	n/a	n/a	n/a	$0.00	
2. Reduce A/B Cost Sharing	$0.00	$0.00	$0.00		$0.00
3. Other A/B Mand Suppl Benefits	0.00	0.00	0.00		0.00
4. Pt B Premium Buydown	0.00	n/a	n/a	0.00	96.40
5. Pt D Premium Buydown Basic	0.00	n/a	n/a	0.00	0.00
6. Pt D Premium Buydown Suppl	0.00	n/a	n/a	0.00	0.00
7. Total	$0.00	$0.00	$0.00	$0.00	
			Unalloc. rebate	$0.00	

Note: The worksheet insurers filled out to allocate their rebate in 2009. Note that Part B premiums are not more difficult to reduce than others.

Table A.1: Example Plan Design

Before rebate allocation:

Coverage	Mandatory	Premium	Example plan
In/outpatient	Yes	96.4	Covered, p=96.4
Supplemental	No	Set by insurer	Drug, podiatry, vision, p=92

After rebate allocation:

Coverage	Mandatory	Premium	Example plan
In/outpatient	Yes	$96.4-Part B rebate	Covered, p=(96-0)=$96.4 AND $47 less cost-sharing
Supplemental	No	Set by insurer, -supplemental rebate	Drug, podiatry, vision, p=(92-23)=$69

Note: This table shows the two layers of MA plan design and how the rebate may be used to change plan features. Our example plan has a $70 rebate, which it uses to reduce $0 of the Part B premium, $23 of the supplemental premium, and $47 of cost-sharing requirements.

Table A.2: Variation Within Contract-Counties

	Mean	Median	P75	P95	P99
Spread rebate	28.77	23.31	47.32	879.56	116.96
IQR rebate	18.04	15.53	29.34	49.69	73.05
Spread supp. premium	35.53	20.00	59.00	127.9	194.00
IQR supp. premium	22.20	11	34.90	87	110.25
Spread cost sharing rebate	25.03	18.35	43.11	74.17	89.95
IQR cost sharing rebate	15.87	15.85	27.11	44.98	61.29
Spread supp. premium rebate	10.42	7	16.73	33.92	57.69
IQR supp. premium rebate	6.55	3.97	10.77	22.96	34.82
Spread Part B rebate	2.47	0.00	0.00	10	93.50
IQR Part B rebate	1.32	0.00	0.00	0.00	51.75
N	20,505				

Note: An observation is a contract-county over the course of the sample period. Spread X is the range of X observed on plans within that contract-county during the sample period. Similarly, IQR X is the interquartile range of X.

Table A.3: Demand Estimates with OS and Part D Rebate Split

D.V. = log market share ratio	Low-Risk β (S.E.)	Medium-Risk β (S.E.)	High-Risk β (S.E.)	Aggregate β (S.E.)
O.S. rebate	1.05	.76	.52	.77
	(.85)	(.87)	(.9)	(.85)
O.S. rebate ^{2}	.031	.04	.042	.04
	(.028)	(.03)	(.031)	(.029)
Part D rebate	2.6***	2.61***	2.46***	2.54***
	(.52)	(.6)	(.61)	(.56)
Part D rebate^2	.012	.0097	.0085	.0093
	(.0089)	(.0095)	(.0096)	(.0091)
Cost-sharing rebate	1.19***	1.31***	1.44***	1.31***
	(.28)	(.3)	(.34)	(.3)
Cost-sharing rebate^2	-.0051**	-.0046*	.0041	-.0042*
	(.0023)	(.0025)	(.0027)	(.0025)
Part B rebate	.34	.0088	.31	.068
	(.37)	(.41)	(.46)	(.41)
Part D bundled	95.8***	82.9***	75.1***	82.1***
	(24.1)	(25.7)	(26.9)	(25.2)
I(Supp. coverage)	-27.7**	-28.3**	-29.2**	-28.2**
	(11)	(11.9)	(13.4)	(12)
Total Part D premium	-1.62**	-1.3*	1.02	-1.29*
	(.71)	(.74)	(.76)	(.71)
Total Part D premium^2	.0073	.0065	.0075	.0077
	(.0065)	(.0066)	(.0066)	(.0064)
Total OS premium	.06	.63*	1.15***	.71**
	(.3)	(.32)	(.34)	(.31)
Total OS premium^2	.00047	.000087	-.00055	.000016
	(.0015)	(.0017)	(.0017)	(.0016)
Constant	-455.6***	-480.9***	-509.9***	-483.2***
	(17.5)	(18.9)	(19.3)	(17.8)
N	90,061	89,187	88,066	90,070

$^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$

Note: Here we split the OS rebate and Part D rebate variables, though our model predicts they should enter utility identically. Traditional Medicare is the outside option. All rebate variables and premiums are taken in dollars. The supplemental premium rebate is the sum of the rebates directed to the Part D premium and the OS premium. The specification also includes an indicator for bundled Part D coverage or other supplemental packages, as as controls for before-rebate Part D coverage or other supplemental packages, as well as controls for before-rebate Pard D and OS premiums. We winsorize Total OS Premium. Observations change across risk segments because we censor plans with less than 10 beneficiaries in each segment.