Low Income Subsidies and Automatic Enrollment in Medicare Part D

Daniel P. Miller
Clemson University *

April 30, 2013

Abstract

In this paper, I examine the design of the Low Income Subsidy (LIS) program that is part of the Medicare Part D prescription drug program. Two key features of the program, subsidy payments and an automatic enrollment provision distort insurers' plan offerings and pricing policies. I develop a structural equilibrium supply/demand model of the market and estimate the model using data on all plan offerings from 2006 to 2009. I conduct two counterfactuals to test whether they reduce the distortions of the LIS program. First, I recompute the market equilibrium under a policy that removes LIS enrollees from the market to see what the market outcomes would be if only regular (non low income) enrollees were served by the Part D market. Second, I preserve the LIS program, but alter the formula for benchmarking subsidy payments and plans' eligibility for auto enrollment to be tied to the enrollment patterns of Non-LIS enrollment. This reduces the externalities of the LIS program on regular enrollees, while having minimal impact on LIS enrollees.

1 Introduction

Following the passage of the 2010 "Patient Protection and Affordable Care Act" federal and state government agencies have been working to design health insurance exchanges that should be ready to go online in 2014. One of the challenges to developing these exchanges is to ensure that they can serve a variety of socioeconomic populations. The Medicare Part D prescription drug program—that has been operational since 2006—is a working example of a health insurance exchange designed to serve a broad base of enrollees. In this paper, I focus on the design, and impact on the market, of the Low Income Subsidy (LIS) program that

*dmille7@clemson.edu

provides additional subsidies for eligible low income Medicare beneficiaries. They compose about 20% of the Medicare population and 45% of the enrollees in Part D drug plans.

The Medicare Part D prescription drug program established a regulated and subsidized health insurance exchange where Medicare beneficiaries can select from a menu of prescription drug plans offered by competing private insurers. Prior to Part D, the typical (non low income) Medicare beneficiary, had no drug coverage under Medicare or may have been able to access supplemental coverage through some other channel. However, many low income Medicare beneficiaries qualify for Medicaid. Before Part D, these "dual eligible" individuals had drug coverage through Medicaid's drug program. After the Part D exchanges became operational, these dual eligible individuals lost their Medicaid drug coverage and were moved over to coverage being offered by private Part D insurers.

The designers of the Part D program aimed to create a market mechanism based on principles of free market exchange and competition amongst insurers. For regular, non low income enrollees they select plans based on price and coverage features, which in turn gives plans incentives to compete. For low income enrollees, it was necessary to offer them affordable, in fact free drug coverage, and to furthermore guarantee that each person is covered even if that means the government needs to make the plan assignment decision on behalf of the beneficiary. The challenge for the designers of the program was coming up with a market mechanism to accommodate both segments of the market (regular and low income) in a way that preserves competition incentives.

The Part D legislations has special stipulations for low income and Medicaid beneficiaries in the Low Income Subsidy (LIS) program. There are two key features. First, the LIS program provides additional government subsidies for monthly premiums and copays over and above those available to regular (non low income) enrollees. For the lowest income levels, drug coverage is free. Second, Medicare/Medicaid dual eligibles, who do not actively choose a plan are automatically and randomly assigned to an LIS eligible plan. To be LIS eligible and hence eligible to receive auto-enrollees, plans must price below a market determined threshold. From year-to-year plans that maintain their eligibility have their auto-enrollees rolled over, and plans that lose their eligibility lose their auto-enrollees who are then randomly assigned to another plan.

These two features decouple plan choice from market fundamentals which in distorts insurers' decisions about plan offerings and pricing. First, is the rather obvious effect that the premium subsidies make enrollees less price sensitive, and thus insurers face more inelastic demand which tends to raise prices. The second feature, automatic enrollment has a less obvious effect. On one hand, when enrollees are automatically and randomly assigned to plans, without regards to prices or plan characteristics, their demand becomes perfectly inelastic, which gives plans incentives to raise price. On the other hand, insurers must price their plans below a market determined threshold, which gives them incentives to keep prices low. The threshold creates a large gap in demand. Pricing just below the threshold gives the plan a large boost in enrollment through those that are automatically enrolled. This demand discontinuity gives rise to a large bunching of prices just below the threshold. But the incentives are more complicated. The threshold is not exogenously set by the regulators. Instead it is based on an enrollment weighted average of plan prices. This creates externalities across plans. For example, an insurer may price its non-LIS eligible plan high with the intend of raising the threshold to benefit its LIS plan. Conversely, it may want to price the other plan low to lower the threshold such that its LIS plan is the only one receiving auto-enrollees.

These features of the LIS program have a large effect on the market. Prior to 2013, Part D prices had steadily risen at a rate of about 5% per year. For 2013, the average bid (from which premiums are determined) fell by 6% because two of the larger insurers decided to price their plans below the LIS threshold. However, there can be rather unusual pricing behavior. In 2009, 12% of insurers' enhanced coverage plans (which are ineligible for auto-enrollees due to the fact they offer enhanced coverage) were priced lower than their companies' basic plans that offer less generous coverage.

In this paper, I first model demand for Part D plans in a differentiated product discrete choice framework and take into account the subsidies and automatic enrollment provisions of the LIS program. I then model the supply-side pricing decisions of multi-product insurers, documenting how automatic enrollment induces price bunching at the LIS threshold, and showing how the subsidies distort residual demand elasticities. Key to this modeling framework is showing how the firms' pricing of other plans, such as its enhanced plans that are ineligible for LIS enrollees, are made with the intend of boosting profitability on its LIS plans.

The demand/supply model is written with respect to static (single period) optimization, however the rules on determining subsidies payments and the threshold have a link to past years. In particular the weights used to determine average bids and hence subsidy amounts and the LIS threshold are based on enrollment figures of LIS enrollees from the previous year. This gives rise to interesting dynamic effects, in particular a cycling of plans pricing low then high, in and out, of LIS eligibility.

After writing out the model, and providing descriptive evidence on firm behavior and enrollment patterns, I estimate the model using the methods of Berry (1994) and Berry Levinsohn Pakes (1995). The standard First Order Condition approach for estimating the supply-side of the model does not apply because of the discontinuity in demand at the LIS threshold. To circumvent this issue, I apply different inference techniques to back out cost.

After estimating the model parameters, I perform two counterfactual exercises. The first removes LIS individuals from the market. This reflects proposals that have been made to return dual eligibles back to Medicaid coverage. Under this counterfactual, the distorted pricing incentives of the LIS program and it serves as a gauge to quantify the distortions in welfare terms. The second counterfactual, preserves the two features of the LIS program, but uses a different formula for calculating LIS subsidy payments and the benchmark. Specifically, I re-weight the formula to be tied to non-LIS enrollment as opposed to the current mechanism that ties it to LIS enrollment. This has the potential to boost competitive incentives in the sense that non-LIS weights reflect the decision making process of enrollees who decide based on price and product quality, whereas the LIS weights are largely driven by randomization of auto-enrollees.

The remainder of paper the follows. In section 2 I review the related literature. In section I describe the background on Medicare Part D. Section 3 presents the demand and supply model. Section 4 provides descriptive evidence and describes the data. Section 5 describes the estimation technique. Section 6 presents the counterfactual results, and 7 concludes.

2 Contribution to Existing Literature

This paper contributes to a literature that evaluates the performance of the Medicare Part D prescription drug program and more broadly government designed health insurance programs. The related literature falls into four broad strands: insurer competition, adverse selection and risk adjustments, consumer choice, and dynamic effects.

Most closely related to this paper is work on the competitive effects of the Part D market. Lucarelli et al. (2011); Miller and Yeo (2012a) develop equilibrium models of the Part D market and simulate counterfactual policies that restrict plans offerings in the former and add a government run public option in the latter. Similar approaches to study competition have been conducted in the Medicare Part C market Town and Liu (2003). This paper builds on the equilibrium modeling approach to assess competitive outcomes. Whereas the other approaches absorbed the LIS program feature in a reduced form manner to the demand models, this paper contributes by explicitly model and estimating the competitive effects of the LIS program.

In the most general context, government insurance programs involve some form of subsidy payments which can distort market outcomes. A well studied distortion comes from risk adjustments designed to correct for adverse selection. Under risk adjustments, plans receive additional payments from the government if they attract a high cost risk pool of enrollees based on their demographic and health status. As shown in Brown et al. (2011) for the Medicare Part C, risk adjustment schemes do not always work well. Plans are able to target low cost enrollees (conditional on their risk profiles), and reforms that have implemented finer calibrations of risk adjustors make matters worse. In a structural model of competition and adverse selection in the Part C market, Lustig (2010) identifies a high degree of adverse selection. While the issues surrounding adverse selection and risk adjustments are likely present in the Part D market, Part D is unique in that a large fraction of the market is randomly assigned to plans under the LIS program. This feature of the program raises its own unique issues regarding plan strategies and market distortions, which dollar-for-dollar are just as important as adverse selection. In fact, adverse selection concerns may be mitigated by random assignment due to the fact enrollees are choosing plans based on lottery outcomes, rather than plan features designed to select low cost enrollees. I take up this idea in concluding remarks to give evidence of whether the tools that plans use to select low cost enrollees are or are not being used by LIS eligible plans.

Another branch of the Part D literature uses individual-level data to study the decision making process of how consumers choose plans. The work in Heiss, McFadden, and Winter (2006, 2010); Abaluck and Gruber (2011); Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2012); Abaluck and Gruber (2011); Ketcham, Lucarelli, Miravete, and Roebuck (2011). documents evidence that they are not choosing plans optimally. The reasons could be that these are complex financial contracts, and it is difficult for senior citizens to sift through documentation on the many plans being offered to select the best one. They may also be biased into selecting plans based on easily observable plan characteristics such as the premium and deductible, while the details of individual drug coverage are more difficult to evaluate. This literature does not consider automatic enrollees due to the fact they are not choosing their own plans. But the issue of whether a plan is an appropriate fit matter for auto-enrollees. Despite facing zero premiums and copays, the auto-enrolled individual must still conform to the plan's drug formulary and restrictions, which may be sub-optimal. Moreover, year-to-year reassignment of LIS enrollees whose plans lose eligibility can impose costs an individual who has conformed to a specific plan.

Another important literature considers the dynamic features of health insurance markets with respect to switching costs. From individual level-data **?** have documented strong persistence in the demand for health insurance, which likely stems from switching costs. Handel (2013); Nosal (2012) estimate switching cost models for the Medicare Part C plans and in an employer-sponsored health insurance setting. Both find very large switching costs on the order of $2000 to $4000. Miller and Yeo (2012b) estimates a dynamic switching cost model for Medicare Part D and like the others finds large switching costs. However, the switching cost matters are more subtle in Part D because of the LIS program. Approximately 3/4 of all switches involve reassignment of LIS enrollees, and in fact there may be "staying costs" for enrollees who face reassignment that wish to stick with their current plan. The identification strategy in their dynamic model of Part D uses year-to-year variation in plan prices and characteristics, which as this paper documents can in part be attributed to a cycling pattern of plans in and out of LIS eligibility. In environments with switching costs, firm strategies take on a dynamic component whereby plans either price low to invest in a stock of new enrollees or price high to harvest profits from locked-in enrollees **?**. This paper shows another reason why a low-then-high pricing can be induced by the program design.

3 Medicare Background

Medicare is the United State's entitlement program that provides health insurance to all people over age 65 and to some categories of disabled people. It started in 1965 and is funded by a payroll tax. The original Medicare programs (Part A and Part B) cover hospital and doctor services, but lack coverage for prescription drugs. Under Part A and Part B, there is only 1 insurance plan that is provided by the government.

Medicare reforms have introduced a privatization of insurance. In 1997, Part C, currently called Medicare Advantage (MA), created a health insurance exchange that gave Medicare beneficiaries the option to purchase plans offered by competing private insurers in lieu of Part A and B coverage. The government regulates coverage and subsidizes MA plans using a scheme that provides subsidy amounts on par with those going to Part A and B, with risk adjustments made to account for plans attracting relatively high or low cost pools of enrollees.

As part of the legislation in the 2003 Medicare Modernization Act (MMA), Medicare introduced Part D to provide prescription drug coverage. All Medicare beneficiaries must obtain coverage or face a penalty. Like Part C, Medicare beneficiaries can choose from a menu of plans sold on an insurance exchange by competing private insurers, however there is no government option available. Consumers can only get coverage from private sources including the plans sold on the Part D exchange or plans sponsored by employer/retiree benefits packages. Notably, the approximately 10 million Medicare beneficiaries who previously obtained drug coverage through the Medicaid program were moved to Part D. Medicaid is an insurance program for low income households of all age ranges, not just Medicare ages. It provides health insurance and includes long term care coverage which is particularly relevant for the Medicaid population that is also eligible for Medicare health coverage.

Medicare regulates drugs coverage by setting a minimum standard. The 2006 minimum standard includes a $250 deductible, a 25% coinsurance rate on drug expenditures up to a limit of $2250 in drug expenditures, then their is a gap in coverage up to $5200, the so called donut hole, and coverage receives at a 5% coinsurance rate for expenditures thereafter. All insurers participating in the market must offer at least one "basic" plan that meets or is actuarially equivalent to the minimum standard. They may also offer "enhanced" plans with more generous coverage, that typically reduce the deductible, provide coverage in the donut hole gap, and/or lower coinsurance rates in the initial coverage zone. Plans are not required to cover the entire universe of prescription drugs on their drug formularies (list of covered drugs). They have a lot of flexibility in deciding what drugs to cover and at what prices. They may exclude drugs, so long as the plan includes some drugs in all of the major therapeutic classes. They can use formulary management techniques including tiered copay/coinsurance rates to differentially price drugs, and place usage restrictions on drugs, such as requirements for prior authorization.

Each year in June, insurers submit proposals for plans they intend to offer in the upcoming year. Markets are geographically separated along state boundaries into 39 regions dividing the US and its territories. Insurers submit separate plans for each market they enter. Along with the proposals, insurers submit bids that serve as the basis for pricing and payments to plans. Medicare subsidizes approximately 2/3 of average bid submitted by plans and the enrollees pay the remainder as a monthly premium. The base subsidy applies for all enrollees. The bid is normalized by each plans risk adjustment factor such that it would represent the cost of enrolling the average Medicare beneficiary. For example if a plan expects to attract a pool of enrollees with health and demographic factors indicating 10% higher cost, the risk adjustors adjust the bid down by 10% to be on par with the cost for an average enrollee.

Low income enrollees have special rules under the LIS program. They account for about 20% of the Medicare population and 45% of all enrollees in Part D plans. They receive additional subsidies for premiums and copays. Households below 100% of the poverty level pay zero premiums and no copays/coinsurance if they select a plan the qualifies for the full low income subsidy. Slightly less impoverished households receive partial subsidies for premiums and copays. The subsidy gradually phases out up until 150% of the poverty level. Low income subsidy enrollees may choose any plan, but they are only subsidized for the portion of the premium attributable to basic coverage, supplemental coverage from an enhanced plan is not subsidized. The enrollee must pay the difference in cost for a basic plan priced above a market determined subsidy threshold.

The second important feature of the LIS program is automatic enrollment. All households that are "dual eligibles," meaning that they are covered under Medicare and Medicaid, will be automatically and randomly assigned to a plan if they do not actively enroll in a plan. They can only be assigned to plans that qualify for the full low income subsidy. That is, basic plans priced below the market determined threshold. Dual eligibles, regardless of where they stand in relation to poverty guidelines qualify for the 100% subsidy. After being notified of assignment, auto-enrollees have ample time to disenroll from the plan and actively select another plan. In subsequent years, an auto-enrollee will be re-enrolled in the same plan if they do not actively choose another plan and if the plan retains its full low income subsidy eligibility. If the plan loses its eligibility, the auto-enrollee is randomly assigned to another qualifying plan. If at any time a dual eligible actively selects a plan, he will never be subject to random plan assignment. A concern that Medicare had with the automatic enrollment provision was the reassignment due to plans losing eligibility occurred frequently. For 2009, Medicare implemented a "de minimimus" policy whereby plans that were LIS eligible and priced above the threshold by less than $2 would retain their auto-enrollees.

4 Model

We model the supply and demand system for plans using the discrete choice framework of Berry (1994); Berry et al. (1995). We separately introduce the demand and supply side.

4.1 Demand non-LIS enrollees

Every year t, a consumer, indexed by i, can enroll in one prescription drug plan. Consumers choose amongst the $j = 1, \ldots, J_{mt}$ differentiated plans offered in market m in year t. Markets are geographically separated by Medicare regions drawn around state borders. Residency requirements and the annual enrollment period admit a very clean market definition; consumers cannot enroll in plans outside their region, nor switch plans within a year. They may also choose an outside option, $j = 0$. We normalize the utility from the outside option to zero. The outside option includes foregoing drug coverage, enrolling in an MA+Part D plan, or getting it from another source, such as a current employer, another government program, or a Retiree Drug Subsidy program plan.

Each year, enrollees pay a premium p_{jmt} set by the plan. They derive utility from plan characteristics and income left over after paying the premium. Define the conditional indirect utility of person i choosing plan j in market m as:

$$U_i(X_{jmt}, p_{jmt}) = -\alpha_i p_{jmt} + \mathbf{X}'_{jmt}\beta_i + \xi_{jmt} + \epsilon_{ijmt} \tag{1}$$

where $\mathbf{X}_{jmt}$ is a vector of observable plan characteristics, ξ_{jmt} represents an index of unobservable (to the econometrician) plan characteristics, and ϵ_{ijmt} is a term capturing idiosyncratic differences in consumers' preferences over plans. The terms α_i, and β_i are random coefficients that represent consumer i's marginal utility over income and over product characteristics. The random coefficients are distributed iid normal across consumers and markets with mean $\bar{\alpha}$ and $\bar{\beta}$ and variance Σ. Consumers choose the plan yielding the highest conditional indirect utility in equation 1.

After describing the supply-side model, we introduce our measures of plan characteristics in the data section and further discuss factors driving heterogeneity in preferences.

4.1.1 Demand with Low Income Subsidy

The above utility specification in equation 1 can be explicitly modified to account for the low income subsidy (LIS). The low income subsidy has three features that we model: premium subsidies, drug cost sharing reductions, and a provision in the legislation that automatically enrolls households that do not actively select a plan.

An enrollee's eligibility for the low income subsidy can be described by the parameter $\kappa_i \in [0, 1]$. A enrollee with $\kappa_i = 1$ is eligible for the full subsidy, $\kappa_i = 0$ is a regular enrollee with no eligibility, and enrollees with values in between receive a partial subsidy. Eligibility is determined in three ways. First, all households that are enrolled in Medicaid are automatically granted eligibility at $\kappa_i = 1$. Second, non-Medicaid households can become eligible through a means test of income and wealth indexed to official Federal Poverty Line (FPL) guidelines. For households below 135% of the FPL $\kappa_i = 1$, and in 0.25 intervals it decreases until income is above 150% of the FPL. Third, other factors such as disability and whether the person is under the care of an institution determine eligibility.

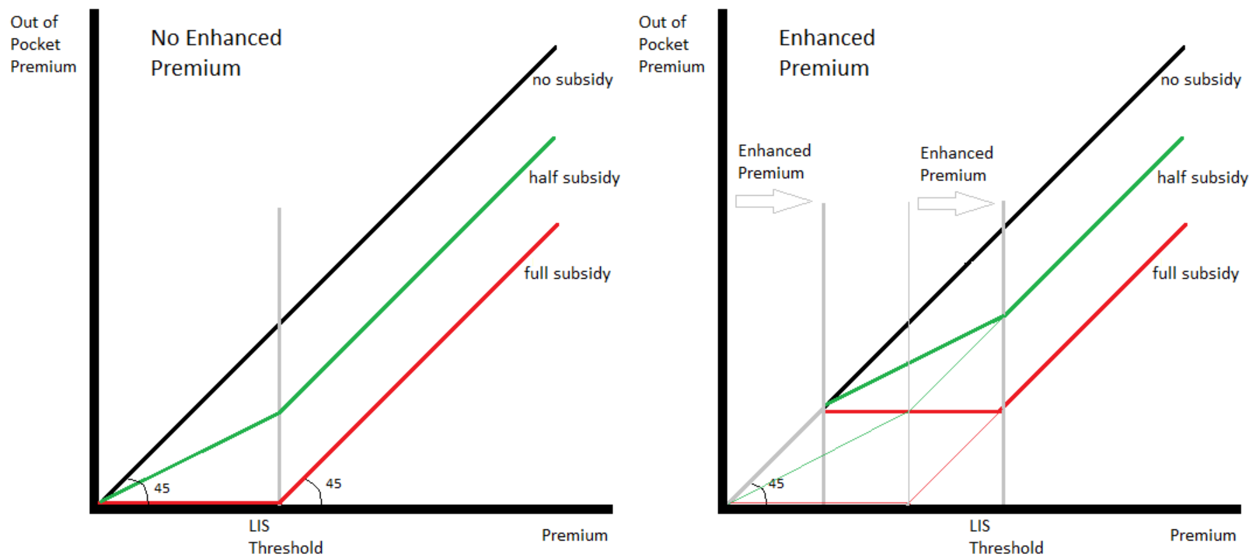

Figure 1: Low Income Subsidy Out-of-Pocket Premium

The low income premium subsidy in market m for plan j, $\tilde{s}_{jmt}$, is a market determined amount. Only the component of the premium attributable to basic coverage, p^{basic}_{jmt}, is subsidized. The subsidy amount is capped at a market determined threshold $\bar{s}_{mt}$. The component of the premium attributable to enhanced coverage $p^{enhanced}_{jmt}$ is not subsidized. The total premium is given by $p_{jmt} = p^{basic}_{jmt} + p^{enhanced}_{jmt}$. In the supply side model, we further elaborate on the threshold and how the two components of the premium are determined. For now, we'll note that the threshold is a weighted average of the premiums set by all plans in a market. The subsidy amount is given by:

$$\tilde{s}_{jmt} = min\{p^{basic}_{jmt}, \bar{s}_{mt}\} \tag{2}$$

The subsidy received by an enrollee of type κ_i is $\kappa_i \tilde{s}_{jmt}$. Figure 1 illustrates the after subsidy, out-of-pocket premium $(p_{jmt} - \kappa_i \tilde{s}_{jmt})$ enrollees of various levels of κ_i pay as a function of the premium. The first panel illustrates a basic plan with no enhanced premium. Regular enrollees ($\kappa_i = 0$) that receive no subsidy pay full price, while enrollees with $\kappa_i > 0$ pay less than full price for premiums up to the threshold, then pay the cost difference for a more expensive plan priced above the threshold. Fully eligible LIS enrollees ($\kappa_i = 1$) pay nothing for basic plans priced below the LIS threshold. The second panel illustrates how an enhanced premium translates out the payment schedule.

The utility specification in equation 1 can be modified to include the LIS subsidy. For an enrollee of type κ_i

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}_{jmt}; \kappa_i) = -\alpha_i(p_{jmt} - \kappa_i \tilde{s}_{jmt}) + \mathbf{X}'_{jmt}\beta_i + \xi_{jmt} + \epsilon_{ijmt} \tag{3}$$

Under this specification, an enrollee receives disutility from choosing a high premium plan, yet conditional on the premium and plan characteristics, an enrollee receives more utility from choosing a plan with a higher subsidy. This specification nests the generic utility specification in equation 1 if $\kappa_i = 0$.

We have limited information in the data about enrollee types κ_i. We observe enrollment figures for the subset of the population that has zero low income subsidy eligibility, $\kappa_i = 0$, and for the subset of the population that has some eligibility $\kappa_i > 0$. For the non-eligible segment of market, we impose the restriction that the subsidy amount does not affect utility and use the baseline utility model in equation 1 to estimate demand. For the LIS eligible segment of the market, we include the subsidy amount in utility according to equation 4. Specifically, we estimate the model as

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}_{jmt}; \kappa_i) = -\alpha_i p_{jmt} + \alpha_i^s \tilde{s}_{jmt} + \mathbf{X}'_{jmt}\beta_i + \xi_{jmt} + \epsilon_{ijmt} \tag{4}$$

where $\alpha_i^s = \kappa_i \alpha_i$.

We model both α_i and α_i^s as random coefficients. With limited data on κ_i, we must make assumptions about the joint distribution of α_i and α_i^s to estimate the model for the subset of LIS eligible households. We begin by assuming the distribution of α_i is normal: $\alpha_i \sim N(\bar{\alpha}, \sigma^2)$. There are 3 points to make about the distribution of α_i^s. The first, most restrictive assumption we could make assumes $\kappa_i = \kappa$ for all i. Then α_i^s is perfectly correlated with α_i, and its distribution is scaled proportionally according to κ: $\alpha_i^s \sim N(\kappa\bar{\alpha}, \kappa^2\sigma^2)$. The second assumption we could make does not restrict κ_i to be fixed but assumes α_i and κ_i are independent. Then $\alpha_i^s \sim F_\kappa(N(\kappa\bar{\alpha}, \kappa^2\sigma^2))$, where F_κ is the marginal distribution of κ_i. The third consideration is to relax independence. We would expect positive correlation between α_i and α_i^s if there is diminishing marginal utility of income over the relevant income range which is below 150% of the federal poverty line.

Given these three points, we estimate the model by assuming α_i and α_i^s are distributed multivariate normal with unrestricted means, variances, and correlation coefficient. The unrestricted nature of this parameterization flexibly allows for heterogeneity in κ_i and non-independence. Moreover, it is convenient for estimation because there are only 3 non-linear parameters. But it is still somewhat parameterized in the sense that F_{κ_i} may not induce a normal distribution over α_i^s. In practice, we do not expect this to be too heavy of a restriction because a very large mass of households have $\kappa_i = 1$. Census poverty statistics indicate that 80% of the 65+ aged population with $\kappa_i > 0$ (income¡150% FPL) have $\kappa_i = 1$ (income¡135% FPL). As an approximation, restricting $\alpha_i = \alpha_i^s$ would not be too bad of an assumption. The multivariate normal distribution allows for some departure from this stricter assumption.

The low income subsidy also reduces the deductible amount and copays/coinsurance rates in the initial coverage zone and donut hole. The rules are analogous to those for the premium. The amount that they are reduced depends on κ_i and the threshold amount is given by minimum coverage standards. The maximum deductible is \$0 for $\kappa_i = 1$, \$60 for $\kappa_i \in (0, 1)$ and \$295 for $\kappa_i = 1$. The actuarially equivalent coinsurance rate is capped in the initial coverage zone to 15% for $\kappa_i \in (0, 1)$, 25% for $\kappa_i = 0$, and set to a nominal copay of \$2.40/\$6.00 (generic/branded) for $\kappa_i = 1$. The donut hole is eliminated for all enrollees with $\kappa_i > 0$. It should be noted these are maximum amounts. Enrollees may purchase enhanced plans with lower values. For example, most enhanced plans have a \$0 deductible. For estimation we treat the deductible as a random coefficient that may be correlated with the premium and subsidy amount. We do not include a random coefficient on the initial coverage zone drug prices because it was difficult to pick up an effect using aggregate market share data. We can use the elimination of the donut hole as an exclusion restriction to instrument for the premium.

Finally, we must account for automatic enrollment of low income subsidy enrollees. For $\kappa_i = 1$ households that do not actively enroll in a plan, Medicare randomly assigns them to a plans. Only basic plans with premiums set below the LIS threshold, $\bar{s}_{mt}$, are eligible to receive randomly assigned enrollees. Medicare distributes them uniformly across insurance companies. It is worth noting that they are not forced into random assignment. A household is allowed to choose some other plan beforehand, and there is ample time during the enrollment period for it to choose another option if it doesn't like its random assignment.

From the perspective of a $\kappa_i = 1$ household, neglecting to enroll and accepting random assignment is not such a bad option because it will pay zero premiums, have zero deductibles, and near zero copays/coinsurance. At first glance, they should be indifferent amongst all plans. But consider the utility of such a household,

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}_{jmt}; \kappa_i = 1) = \xi_{jmt} + \epsilon_{ijmt} \tag{5}$$

there are still two terms in the utility function: unobserved plan characteristics, ξ_{jmt}, and the idiosyncratic preference shock, ϵ_{ijmt}. The terms capture average and idiosyncratic unobserved characteristics about the composition and restrictions of the plan's drug formulary and they capture non-fiduciary plan qualities such as the plan's marketing activities and customer service characteristics. Search costs can rationalize the behavior of enrollees that accept random assignment. If the search cost exceeds the difference in utility between the randomly assigned plan and best alternative, the enrollee will accept the randomly assigned plan.

We model automatic enrollment by including a dummy variable in the utility specification that indicates whether a plan is eligible to receive randomly assigned enrollees.

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}_{jmt}; \kappa_i) = -\alpha_i p_{jmt} + \alpha_i^s \tilde{s}_{jmt} + \mathbf{X}'_{jmt}\beta_i + \beta^{lis}\mathbf{1}(LISPLAN_{jmt}) + \xi_{jmt} + \epsilon_{ijmt} \quad (6)$$

With the logit parameterization, ϵ_{ijmt} takes on an additional interpretation as the random number generator determining which plan an individual enrollee gets assigned to. Holding all else fixed in the utility specification, the markets shares s_{jmt} induced by the logit model,

$$s_{jmt} = \frac{exp(\beta^{lis}\mathbf{1}(LISPLAN_{jmt}))}{\sum_k \exp(\beta^{lis}\mathbf{1}(LISPLAN_{kmt}))}$$

distributes automatic enrollees uniformly across LIS plans, and the coefficient β^{lis} determines the proportion of households that are randomly assigned.

There is also a dynamic component related to automatic enrollment. From year to year, automatic enrollees are kept in the same plan if it maintains its qualification. If the plan loses its qualifications, the carry over enrollees are redistributed amongst other plans. We take up this issue further in our switching cost paper Miller and Yeo (2012b), but reassignment will not bias our results in this static model because the aggregate allocation of enrollees remains uniform across LIS eligible insurers.

4.2 Supply

We model the supply side by closely following the regulations in the Medicare Modernization Act of 2003. A set of F multiproduct insurers compete in bertrand nash fashion. In year t, each plan j offered in market m submits a bid b_{jmt} to Medicare. Insurers submit separate bids in each market, even if the plans offered in different markets are otherwise similar. For each enrollee, the plan receives a monthly payment equal to its bid. Part of that payment is made by enrollees in the form of the premium p_{jmt}, and the remainder is subsidized by the government.

We model plans' marginal costs mc_{jmt} of enrolling a customer as a constant. The marginal cost can be separated into the basic and enhanced component.

$$mc_{jmt} = mc_{jmt}^{basic} + mc_{jmt}^{enhanced}$$

where we define the ratio

$$\gamma_{jmt} = \frac{mc_{jmt}^{enhanced}}{mc_{jmt}^{basic}}. \tag{7}$$

For example, γ_{jmt} is zero for basic plans and is larger for an enhanced plan that eliminates the deductible and provides donut hole coverage than for an enhanced plan that only eliminates the deductible.

As multiproduct firms that can offer multiple plans in many regions, profits for firm f are given by,

$$\Pi_{ft} = \sum_{mt} M_{mt} \sum_{J_{fmt}} (b_{jmt} - mc_{jmt}) s_{jmt}(\mathbf{b}) \tag{8}$$

where M_{mt} is the number of potential enrollees in market m and J_{fmt} indexes the set of plans offered by firm f in market m. Notice the market share s_{jmt} is written to explicitly depend on the bid vector $\mathbf{b}$ for all plans across all markets.

The standard approach to solve for marginal cost is to invert a system of first order conditions with respect to bids. In the Bertrand Nash equilibrium, marginal cost is a function of observed bids, markets shares, and a matrix of estimated own and cross price elasticities with respect to bids. There are three complications with this approach. First, the demand system is estimated in terms of premiums, not bids. We need to account for the subsidy rules that map bids into premiums. Second, the demand system exhibits discontinuities because of LIS households that are automatically enrolled in plans with premiums below the LIS threshold. Because of the discontinuities, first order conditions cannot be directly applied. The third complication is adverse selection. Adverse selection is a potential problem for identifying marginal cost because marginal cost would be a function of the bid vector $\mathbf{b}$, not a constant. We consider all of these issues in our identification strategy.

4.3 Premium and Low Income Subsidy Rules

The regulations set the rules for determining the size of the overall premium subsidy, which applies for all enrollees, and the low income subsidy which only applies for LIS enrollees.

First consider the overall premium subsidy. The government subsidizes a fixed proportion, λ_t, of the enrollment weighted *average basic* bid component of all plans in the country ($\lambda_t \approx 65\%$). The enrollee pays the balance as its premium. Thus, each enrollee gets the same subsidy amount regardless of plan choice. Enrollees realize savings from choosing cheaper than average plans, or pay extra to pick one that is more expensive than average. Like the premium, a bid b_{jmt} separates into a basic component b_{jmt}^{basic} and an enhanced component $b_{jmt}^{enhanced}$, where $b_{jmt} = b_{jmt}^{basic} + b_{jmt}^{enhanced}$.

The formula to map a bid b_{jmt} to a premium p_{jmt} is:

$$p_{jmt} = b_{jmt} - \lambda_t \bar{b}_t \tag{9}$$

The weighted average bid $\bar{b}_t$ is based on bids of all stand-alone part D plans and select MA+part D plans in the entire nation.

$$\bar{b}_t = \sum_{jt} \tilde{w}_{jt-1} b_{jt}^{basic} \tag{10}$$

The weights $\tilde{w}_{jt-1}$ are based on the previous year's enrollment E_{jt-1},

$$\tilde{w}_{jt-1} = \frac{E_{jt-1}}{\sum_{jt} E_{jt-1}}$$

The weight is zero for plans that are new entrants to the market. In the first year, 2006, the weights were equal for all plans. The shift from a simple average to the weighted average method was gradually phased in through 2008.[1]

The formulas to separate the basic and enhanced component of the premium are:

$$\begin{aligned} p_{jmt} &= p_{jmt}^{enhanced} + p_{jmt}^{basic} && (11) \\ p_{jmt}^{enhanced} &= b_{jmt}^{enhanced} && (12) \\ p_{jmt}^{basic} &= b_{jmt}^{basic} - \lambda_t \bar{b}_t && (13) \end{aligned}$$

The basic component of the premium is bounded below at zero. It has never been a strictly binding constraint, but, as we will argue in the next section, may have modestly affected bidding strategies for a small number of plans. We choose not to model the constraint.

[1] The "Medicare Demonstration to Limit Annual Changes in Part D Premiums Due to Beneficiary Choice of Low-Cost Plans" act, passed in mid-2006, amended the original legislation to phase-in the weighted average bid calculation method.

In our baseline model, we assume plans choose the bid b_{jmt} but do not have discretion allocating between the basic and enhanced components. The regulations state that the proportion of the bid allocated to each component is based on an actuarial cost calculation that takes into consideration the plan's coverage characteristics. Because the proportion is based on cost, the ratio between the basic and enhanced component of the bid is that same as that between the cost components:

$$\gamma_{jmt} := \frac{mc_{jmt}^{enhanced}}{mc_{jmt}^{basic}} = \frac{b_{jmt}^{enhanced}}{b_{jmt}^{basic}} \tag{14}$$

Later, we relax this assumption by allowing firms some discretion over how to allocate amongst the basic and enhanced components of the bid.

The demand section described the rules for the low income subsidy. A key component is the LIS threshold, $\bar{s}_{mt}$ which is determined by the plans' bids. The threshold is the enrollment weighted *average basic* component of the premium for all plans in a region. The weights are based on the previous year's enrollment.

$$\bar{s}_{mt} = \sum_{jmt} \tilde{w}_{jmt-1}^{lis} p_{jmt}^{basic} \tag{15}$$

The weights $\tilde{w}_{jmt-1}^{lis}$ are based on the previous year's enrollment of LIS eligible enrollees who have $\kappa_i > 0$

$$\tilde{w}_{jt-1}^{lis} = \frac{E_{jt-1}^{lis}}{\sum_{jt} E_{jt-1}^{lis}}$$

The weight is zero for plans that are new entrants to the market. The threshold is bounded below by the minimum premium of a plan that only offers basic coverage. In the program's 7 year history, this has only been a binding constraint once, in Nevada in 2009.

4.4 Residual Demand Elasticities with Subsidy Distortion

The demand model is expressed in terms of the premium, but, for the supply side model, it is necessary to express demand elasticities in terms of bids, not premiums. The subsidy rules distort insurers' residual demand elasticities. The share of enrollees of type (α_i, κ_i) that enroll in plan j in region m in year t is given by:

$$s_{jmt} = \frac{M_{jmt}}{1 + \sum_k M_{kmt}}$$

The term M_{jmt} depends on whether the plan's basic premium is above or below the LIS threshold $\bar{s}_{mt}$. By substituting in the subsidy rules given in equations 9 and 15 those terms are given by

Above low income threshold

$$M_{jmt} = exp\left(-\alpha_i\left(b_{jmt} - \lambda_t \sum_m \sum_k \tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}} - \kappa_i\left(\sum_k \tilde{w}^{lis}_{kmt-1}\left(\frac{b_{kmt}}{1+\gamma_{kmt}} - \lambda_t \sum_m \sum_l \left(\tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}}\right)\right)\right)\right) + \mathbf{X}'_{jmt}\beta + \xi_{jmt}\right).$$

Below low income threshold

$$M_{jmt} = exp\left(-\alpha_i\left((1-\kappa_i)\left(b_{jmt} - \lambda_t \sum_k \tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}}\right) + \kappa_i \frac{\gamma_{jmt}}{1+\gamma_{jmt}} b_{jmt}\right) + \mathbf{X}'_{jmt}\beta + \beta^{lis}\mathbf{1}(\gamma_{jmt}=0) + \xi_{jmt}\right).$$

Notice in particular the inclusion of the term β^{lis}. This reflects enrollment of those low income households automatically assigned to the plan. Only basic plans are eligible to receive automatic enrollees: plans with $\gamma_{jmt} = 0$.

For non-low income subsidy enrollees of type $(\alpha_i, \kappa_i = 0)$ the expression simplifies to

$$M_{jmt} = exp\left(-\alpha_i\left(b_{jmt} - \lambda_t \sum_k \tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}}\right) + \mathbf{X}'_{jmt}\beta + \xi_{jmt}\right).$$

There are three relevant price elasticities: own price, cross price with a plan offered in the same market m, and cross price with a plan offered in a different market m'.[2] Cross price elasticities across markets matter because the overall premium subsidies are based on the bids of all plans across the nation. There is a "kink" in the demand curves at the LIS threshold, so we calculate different elasticities for plans priced above and below the threshold. The LIS threshold does not matter for cross price elasticities with plans in other markets because it is determined market-by-market.

Below low income subsidy threshold

$$\begin{aligned}
\eta_{jjmt} &= \frac{\partial s_{jmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{jmt}} = -\alpha_i b_{jmt}\left[(1-s_{jmt}) - \kappa_i\left(1-\frac{\gamma_{jmt}}{1+\gamma_{jmt}}\right)(1-s_{jmt}) + \kappa_i \frac{w^{lis}_{jmt}}{1+\gamma_{jmt}} s^{above}_{mt} - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjmt} &= \frac{\partial s_{kmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{kmt}} = -\alpha_i b_{jmt}\left[-s_{jmt} + \kappa_i\left(1-\frac{\gamma_{jmt}}{1+\gamma_{jmt}}\right)s_{jmt} + \kappa_i \frac{w^{lis}_{jmt}}{1+\gamma_{jmt}} s^{above}_{mt} - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjm't} &= \frac{\partial s_{kmt}}{\partial b_{jm't}}\frac{b_{jm't}}{s_{kmt}} = -\alpha_i b_{jm't}\left[-(1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jm't}}\tilde{w}_{jm't-1}s_{0mt}\right]
\end{aligned} \tag{16}$$

The first terms inside the brackets for the own and cross price elasticities within the same market are standard for the logit model with no subsidy distortions. The second term

[2] Because the weights $\tilde{w}_{jmt-1}$ are based on lagged enrollment, we could also calculate cross price elasticities across time. We don't because our model is static.

reflects the distortion caused by the low income subsidy. Enrollees with $\kappa_i > 0$ pay a fraction of the premium, which makes the own price residual demand more inelastic. Likewise, those enrollees decreased price sensitivities increases cross price elasticities amongst plans in the same market. The third term is a pricing externality that captures the effect of the bid on the LIS threshold and hence the maximum subsidy amount, $\bar{s}_{mt}$. The intuition is that when a plan increases its bid, it raises the maximum subsidy amount $\bar{s}_{mt}$. The term, s_{mt}^{above} is the market share of plans priced above the LIS threshold. This pricing externality makes the own price residual demand more elastic because above threshold plans are more desirable. Cross price become more smaller. Raising the threshold has no effect on the margin for plans priced below the threshold because the subsidy amount is capped by the premium. Note that this effect is significant for plans with a high weight w_{jmt-1}^{lis} in the calculation of the LIS threshold. The final term reflects the distortion caused by the overall premium subsidy. It makes own price elasticities more inelastic and cross price elasticities larger relative to a market with no subsidy. The intuition is that when plan j in market m increases it's bid, the subsidy increases for all plans across the nation. With a larger subsidy, inside goods become more attractive relative to the outside option. Insurers internalize their marginal effect on the subsidy and will have higher markups, more so for large national insurers with high enrollments (hence high weights $\tilde{w}_{jmt-1}$) that offer plans in many markets. Also notice the subsidy distortion would be more severe if the subsidy fraction λ_t were higher or if Medicare subsidized the enhanced component of bids (γ_{jmt}=0 for enhanced plans). Without subsidies the cross price elasticities with plans in different markets would be zero, but it is positive because the subsidy is determined by the bids of all plans in the nation.

Above low income subsidy threshold

$$
\begin{aligned}
\eta_{jjmt} &= \frac{\partial s_{jmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{jmt}} = -\alpha_i b_{jmt}\left[(1-s_{jmt}) - \kappa_i \frac{w_{jmt}^{lis}}{1+\gamma_{jmt}}(1-s_{mt}^{above}) - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjmt} &= \frac{\partial s_{kmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{kmt}} = -\alpha_i b_{jmt}\left[-s_{jmt} - \kappa_i \frac{w_{jmt}^{lis}}{1+\gamma_{jmt}}(1-s_{mt}^{above}) - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjm't} &= \frac{\partial s_{kmt}}{\partial b_{jm't}}\frac{b_{jm't}}{s_{kmt}} = -\alpha_i b_{jm't}\left[-(1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jm't}}\tilde{w}_{jm't-1}s_{0mt}\right]
\end{aligned}
\qquad (17)
$$

For plans that are above the low income subsidy, the first and third terms are the same as plans that are below the subsidy. But, the second term for plans below the threshold is not present. Because the low income subsidy is capped, marginally changes in the bid affect all enrollees the same regardless of their type κ_i. Thus demand elasticities are not directly affected by the low income subsidy fraction. But there is an indirect effect working through the LIS threshold, which is captured in the second term. If a plan increases its bid, it increases the threshold, which increases the low income subsidy amount for its own low income enrollees. Own price elasticities become more inelastic. As already discussed, the same pricing externality with respect to all other plans priced above the threshold s_{mt}^{above} makes demand more elastic. The final term is the pricing externality with respect to the overall premium.

In our results, we quantify the impact of the subsidy rule on markups by comparing our estimated markups to estimates from a model where insurers treat the overall premium subsidy amount as lump sum and treat the low income subsidy threshold as exogenous.

Furthermore, we must account for automatic enrollment which is determined by the bid. Recall, a plan qualifies for automatic enrolles if $p_{jmt}^{basic} \leq \bar{s}_{mt}$ and it has no enhanced component of the bid ($\gamma_{jmt} = 0$). The expression is modified for a plan below the subsidy by including the term β^{lis}.

$$M_{jmt} = exp\left(-\alpha\left(b_{jmt} - \lambda_t \sum_k \tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}}\right) + \alpha\kappa\left(b_{jmt} - \lambda_t \sum_k \tilde{w}_{kmt-1}\frac{b_{kmt}}{1+\gamma_{kmt}} - \frac{1}{1+\gamma_{jmt}}b_{jmt}\right) + \mathbf{X}'_{jmt}\beta + \beta^{lis} + \xi_{jmt}\right).$$

This will give rise to a discontinuity in the plan's residual demand at the subsidy threshold. The above elasticities assumed fixed α_i and κ_i for illustrative purposes. With random coefficients, aggregate demand and aggregated demand elasticities are calculated by integrating across the distribution of the estimated α_i and α_i^s random coefficients.

4.5 Estimating Marginal Cost

The discontinuity in the demand presents a challenge to estimate marginal cost. Figure 2 illustrates residual demand curves, marginal revenue curves, and the pricing patterns for a variety of marginal costs. The first panel represents basic plans where there is a discontinuity at the LIS threshold. The second panel illustrates enhanced plans, where there is no discontinuity, but there is still a "kink" at the threshold.

.

Pricing exhibits bunching around the LIS threshold for basic plans. Plans that have marginal cost in the range between MC1 and MC4 price at the threshold. The discontinuity also generates a gap in pricing above the threshold, which is illustrated by the range between b1 and the threshold bid. Notice there is no bunching or gap for enhanced plans because demand is continuous.

Figure 3 is a histogram of bids (measured as $p_{jmt}^{basic} - \bar{s}_{mt}$) for basic plans, enhanced plans and all plans together. The pattern of bid bunching and the gap is evident in the data. There is a mass of basic plans that set their premium right at, or very near the threshold.

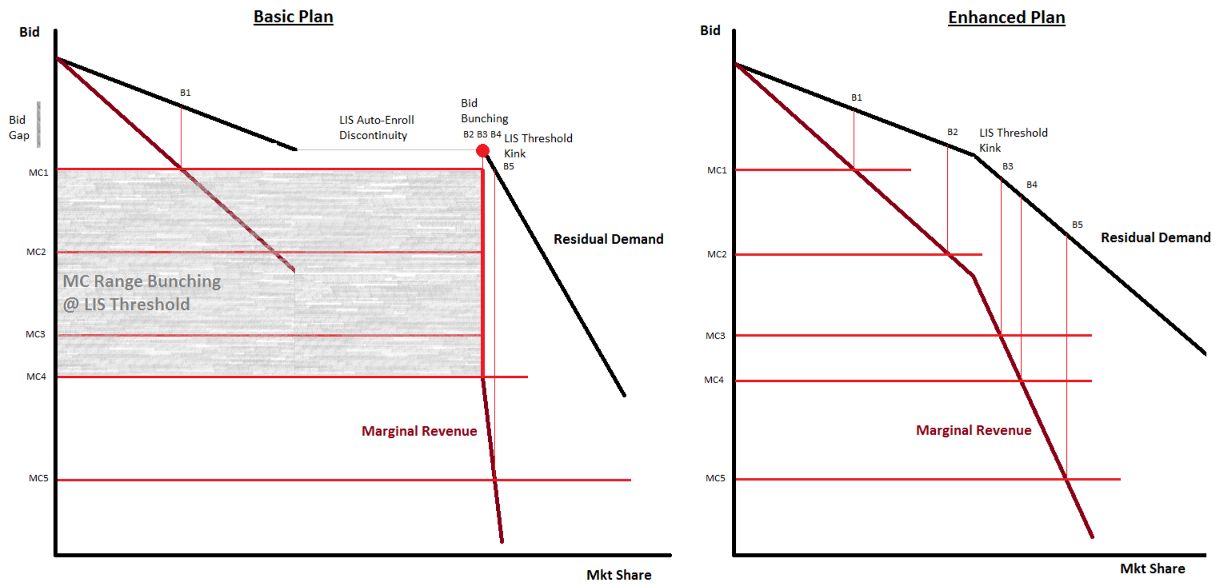

Figure 2: Demand Discontinuity

There is a gap in the number of plans priced just above the threshold, and then a second mode well above the threshold. By comparison, enhanced plans that are not eligible to receive automatically enrolled low income households, do not exhibit any bunching around the threshold.

Because of the bunching at the discontinuity the first order conditions to optimal bidding do not hold with equality. For plans at the threshold there isn't a one-to-one mapping between marginal cost and bids. There exists a range of marginal cost parameters (MC1 through MC4 in figure 2 that would choose to bid at the threshold. The usual procedure of inverting the first order conditions to solve for marginal cost cannot be directly applied.

To circumvent this problem we place a cross-plan restriction on cost. The restriction is about the cost of a basic plan priced at the threshold, and the corresponding enhanced plan offered by the same insurer. The restriction permits us to invert first order conditions to solve for marginal cost.

Recall that the marginal cost of an enhanced plan is additively separable into a basic and enhanced component

$$mc_{jmt} = mc_{jmt}^{basic} + mc_{jmt}^{enhanced}$$

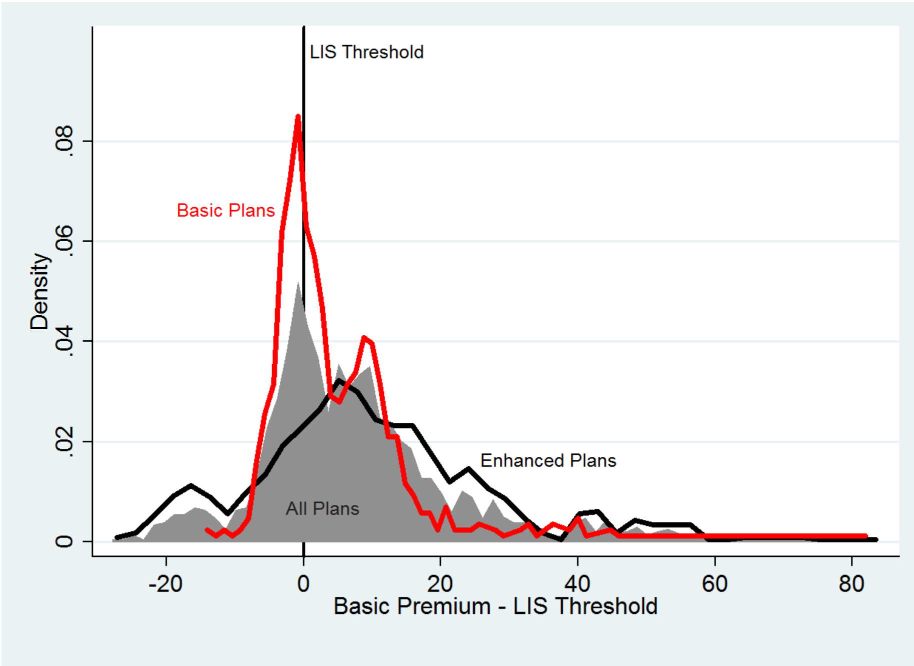

Figure 3: Bid Histogram: Bunching @ LIS Threshold

Similarly, by definition a basic plan has cost

$$mc_{jmt} = mc_{jmt}^{basic} + 0$$

The cross plan restriction states that the basic component of marginal cost on an enhanced plan k equals the marginal cost of that same firm's basic plan j which is offered in the same market.

$$mc_{jmt}^{basic} = mc_{kmt}^{basic} \tag{18}$$

To see how the first order condition can apply consider an example of a firm that offers an enhanced plan (plan 1) and a basic plan (plan 2). The first order condition of equation 8 with respect to the bids is:

$$0 = s_1 + (b_1 - mc_1)\frac{\partial s_1}{\partial b_1} + (b_2 - mc_2)\frac{\partial s_2}{\partial b_1} \tag{19}$$

$$0 \geq s_2 + (b_1 - mc_1)\frac{\partial s_2}{\partial b_1} + (b_2 - mc_2)\frac{\partial s_2}{\partial b_2} \tag{20}$$

If the basic plan 2 is priced away from the low income subsidy threshold, the system of first order conditions can be inverted to solve for both mc_1 and mc_2. If it is priced at the threshold, the second FOC does not hold with equality and the system of equations cannot be inverted. Equality only holds for the first FOC, but the 2 unknowns (mc_1 and mc_2) cannot be solved for because there is just 1 equation.[3] Substituting in the cross plan price restriction and making use of the assumption about the ratio of enhanced and basic marginal cost and bids γ, the FOC for plan 1 becomes

$$0 = s_1 + (b_1 - mc_1)\frac{\partial s_1}{\partial b_1} + \left(b_2 - \frac{mc_1}{1+\gamma_1}\right)\frac{\partial s_2}{\partial b_1} \tag{21}$$

Here, there is one equation and one unknown that can be solved for, mc_1. We can then reapply the restriction to solve for mc_2 using γ.

The restriction can be scaled up for a multiproduct firm serving multiple markets. For a firm that has no plans priced at the threshold the matrix representation of the first order condition is:

$$\mathbf{0} = \mathbf{s} + \Delta(\mathbf{b} - \mathbf{mc}) \tag{22}$$

where the vectors have length N equal to the number of plans offered by the firm across the nation and Δ is the matrix of share derivatives with the jk entry equal to $\frac{\partial s_j}{\partial b_{jk}}$. Note that the FOC cannot be split market-by-market because the subsidy rules create cross-market cross-price elasticities. Marginal cost can be solved for by inverting the system of equations:

$$\mathbf{mc} = \mathbf{b} + \Delta^{-1}\mathbf{s} \tag{23}$$

For firms with plans priced at the threshold the first order conditions are modified by imposing the cost restriction in (18):

$$R\mathbf{0} = R\mathbf{s} + R\Delta\mathbf{b} - R\Delta R'R\mathbf{mc} - R\Delta' M_\gamma R'R\mathbf{mc} \tag{24}$$

The restriction matrix R has dimension $(N - N_{\bar{s}} \times N)$ where $N_{\bar{s}}$ is the number of plans priced at the threshold. The jj entry is a one for the first $j = 1, \ldots N - N_{\bar{s}}$ entries corresponding to plans not priced at the threshold, and the remaining columns are zero vectors which correspond to the plans priced at the threshold. The $(N \times N)$ matrix M_γ indexes the

[3]Strictly speaking, the term $\frac{partials_2}{\partial b_1}$ is only defined for the derivative taken in the negative direction.

enhanced plan and threshold plan for which the cost restriction is imposed. If enhanced plan j is matched with threshold plan k, the jk element takes on the value $1/(1+\gamma_j)$. Note that the γ_j terms are observed in the data because they are equal to the ratio of the enhanced and basic components of the bids.

The restricted system of FOCs can be inverted to solve for the restricted set of marginal costs $R\mathbf{s}$.

$$R\mathbf{mc} = (R\Delta R' + R\Delta' M_\gamma R')^{-1}(R\mathbf{s} + R\Delta\mathbf{b}) \tag{25}$$

The remaining $N_{\bar{s}}$ marginal cost terms can be solved for using knowledge of the γ_j terms and reapplying the cost restriction.

While in theory this approach works, there are few caveats to be aware of when applying the approach. First is the possibility that the inversion matrix does not have full rank. This can occur for two reasons. A few of the firms only offer basic plans and all of them are priced at the threshold. There is little hope in identifying marginal cost. We set marginal cost equal to the average of the other basic plans in the respective region. More generally, full rank fails if for some threshold plan j in market m there does not exists a corresponding enhanced plan k. The regulations require that each firm offering an enhanced plan, must also offer a basic plan. The converse is not true; insurers are not required to offer an enhanced plans. This binds in a few cases; an insurer may offer enhanced plans in many markets, but not offering one in just a few. Again, we set marginal cost of that threshold plan to the average of the other firms' basic plans in the region. This should have minimal effect on the marginal cost estimates for plans in other regions because the cross-market cross plan elasticities are small.

The second issue is about the selection of which enhanced plan should be matched to which basic plan. Many firms offer 2 enhanced and 1 basic plan in a region. We choose the enhanced plan with observed product characteristics closest to the basic plan.

The third, and most substantive issue, regards the possibility of incomplete information. In the histogram of bids in figure 3, many plans do not set their price exactly at the threshold. With incomplete information, the firms would price within a few dollars of the threshold. There are also several plans that appeared as if they intended to price at the threshold, but actually end up pricing just above it. Including incomplete information greatly complicates the model. Instead we designate any plan that prices within a small dollar range ($2) of the threshold as being a threshold plan.

4.6 Allocation of Common Costs

The fourth issue, and perhaps most relevant for the strategic aspects of bidding, is the assumption that insurers have no discretion in allocating bids amongst the basic and enhanced component of the bid. To relax the assumption, we allow insurers discretion in how they allocate profits across the basic and enhanced component of the bid. Formally, we interpret the regulations as stating,

$$mc_{jmt}^{enhanced} \leq b_{jmt}^{enhanced} \tag{26}$$

$$mc_{jmt}^{basic} \leq b_{jmt}^{basic} \tag{27}$$

$$\tag{28}$$

The proportionality assumption on γ_{jmt} implicitly assumes plans allocate profits proportionally across the basic and enhanced components of the bid; the above inequalities allow plans to allocate all of their profits to the basic component, all to the enhanced component, or somewhere in between.

There is reason to believe plans have incentives to allocate at one extreme or the other. The idea is that a firm offering a *basic* plan priced at or below the threshold has an incentive to keep the threshold as low as possible. Because the LIS subsidy amount is capped by the premium, there is no benefit to having a higher LIS threshold. A higher LIS threshold only benefits plans priced above the threshold. The insurer can allocate profits on its *enhanced* plan to the enhanced component of the bid, thereby lowering its basic component of the bid and keeping the LIS threshold low. On the other hand, an insurer offering a *basic* plan priced above threshold benefits from a higher LIS threshold. It can allocate profits on its *enhanced* plan to the basic component of the bid, thereby raising its basic component of the bid and keeping the LIS threshold high. This intuition is present in the elasticity equations through the LIS threshold pricing externality. But note, the ability of the enhanced plan to have any effect on the LIS depends on its lagged enrollment of LIS enrollees. These incentives are strongest for plans with high lagged enrollment.

We explore this incentive in a reduced form manner through the following regression of the enhanced component of the bid $b_{jm}^{enhanced}$ for all enhanced plans as a function of whether or not the insurer f offered a companion basic plan in the same market below the LIS threshold

Table 1: Allocating Enhanced Component of Bid

	(1)	(2)
LISfirm	1.81	-.96
	(.44)	(.41)
LISfirm*lagLISshare	42.7	-77.6
	(26.5)	(21.7)
Deduct	-.71	-1.90
	(.06)	(.18)
PI	-.24	-.06
	(.04)	(.08)
PIdonut	-1.11	-1.05
	(.04)	(.03)
Form100	-.46	.25
	(.06)	(.12)
FE	Y	N

$LISfirm^{\epsilon}_{jm}$: basic plan priced within $3 of LIS threshold avg lagLISshare: .003 lagLISshare¿.01: 7.5%...a few enhanced plans can affect LIS threshold

($LISfirm^{\epsilon}_{jm} = 1$) or above the LIS threshold ($LISfirm^{\epsilon}_{jm} = 0$)in the same market.

$$b^{enh}_{jm} = \mathbf{X}'_{jm}\beta + \beta LISfirm^{\epsilon}_{jm} + \beta LISfirm^{\epsilon}_{jm} lagLISshare_{jm} + fe_m + fe_f + \epsilon_{jm}$$

We included extensive controls for plan characteristics, market fixed effects and insurer fixed effects. We also interact the LISfirm dummy variable with the lagged LIS share for enhanced plan jm because the theory predicts plans with larger lagged shares have a greater effect on the threshold. Note basic plans are excluded from the regression because, by definition, the enhanced component of the bid is $0 and cannot be strategically allocated.

The results in column 2 of table 1, that excludes firm fixed effects, shows that the enhanced component of the bid for LIS firms are lower than that for non LIS firms. This could simply reflect the fact that there is correlation in the cost of basic and enhanced plans within a firm. In other words, low cost firms are more likely to be LIS firms, and as such will have lower costs for their basic plans, which in turn is reflected in a lower enhanced bid. The first column controls for correlated firm costs by including firm fixed effects. These results show that the enhanced component of the bid is lower for LIS firms than non-LIS firms. This is consistent with the prediction that firms strategically allocate bids on the enhanced component. Absent the strategic effect, the coefficient estimate would be zero. The second line in the table shows the interaction effect of LIS firm with the lagged LIS share of the enhanced plan. Consistent with strategic allocation, firms with larger shares allocate more to the enhanced bid. Given the average enhanced plan has a small share of LIS enrollees, this effect is quite tiny in dollar values, but some plans, 7.5%, have a lagged share greater than .01. For these plans the dollar term translates to an additional $0.42 increase in strategic allocation.

4.7 Dynamics: Plans cycling in and out of LIS eligibility

The supply-side model has an inherent dynamic component because the determination of the overall subsidy and LIS threshold depends on lagged enrollment. For the LIS threshold, this creates an interesting dynamic effect. From the FOC, plans that have a large share of lagged LIS enrollees are more likely to price above the threshold, and plans that either have no lagged LIS enrollees, or a small share, are more likely to price at the LIS threshold. Given the majority of LIS enrollees are assigned through automatic enrollment, LIS plans in year $t-1$ are less likely to be LIS plans in year t. Likewise, non-LIS plans in $t-1$ are more likely to be LIS in year t.

These incentives can create a cycling effect of plans in and out of LIS eligibility. This of great concern to policy makers, because LIS enrollees get moved in and out of plans each year. Reassignment imposes switching costs on enrollees, and there is a desire to keep their coverage stable from year to year. The "de minimimus" policy on plan LIS eligibility was put in place with the specific aim of reducing year-to-year reassignment.

In preliminary analysis we present two pieces of evidence to support this dynamic effect. The histogram in figure 4 is similar to the one in figure 3, except it restricts the sample to plans that are new entrants to the market. By definition new entrant plans have zero lagged LIS enrollees and thus have no effect on the LIS threshold. The figure shows the same pattern of bunching at the threshold, but what is different is that there is a higher fraction of plans pricing at the threshold. In other words, new plans are more likely price as LIS plans. While this pattern could be due to some other reason, such as new plans having lower cost, it is suggestive that lagged LIS enrollment matters for pricing with respect to the LIS threshold.

The second piece of evidence in table 2 considers the time series serial correlation of plan characteristics and LIS eligibility.

The second column reports the AR(1) correlation between year $t-1$ and year t plan characteristics. For key variables such the premium, deductible, and basic/enhanced status, there is a high degree of correlation around 0.9. It is less for the coverage characteristics

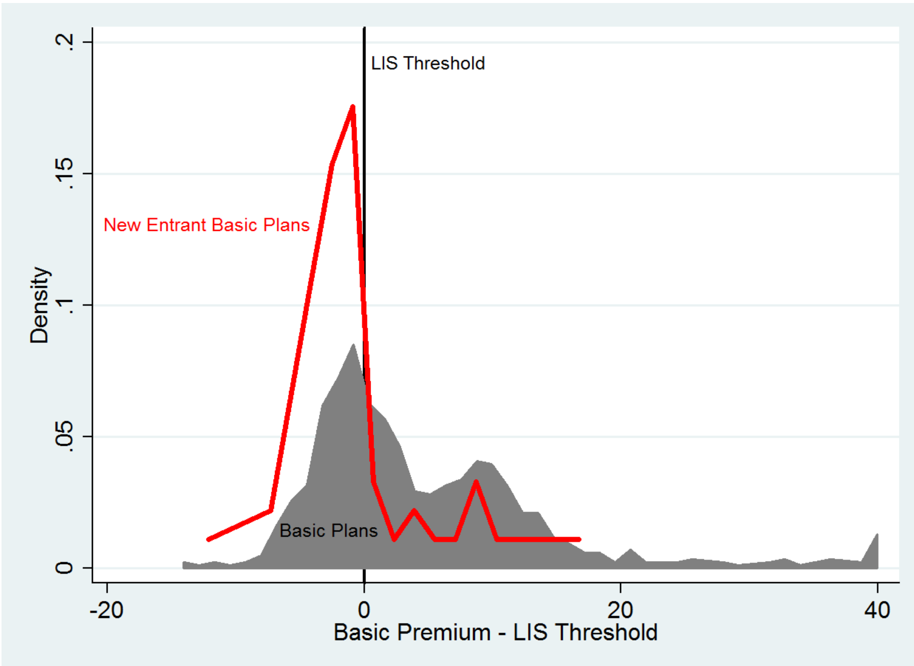

Figure 4: Bid Histogram: New Entrant Bunching @ LIS Threshold

Table 2: AR(1) Process of Product Characteristics

X	Constant	AR(1) coefficient	R^2
Monthly Premium	4.80*	0.95*	0.65
Monthly Deductible	0.77*	0.87*	0.73
Benefit EA dummy	0.09*	0.90*	0.81
Initial Coverage PI	12.69*	0.71*	0.49
Donut Hole PI	19.79*	0.79*	0.64
pharm per eligible (x1000)	0.29*	0.78*	0.78
LIS dummy	0.06*	0.60*	0.42

The regression equation is $X_t = \gamma_0 + \gamma_1 X_{t-1}$. The results are based on 4270 observations across years and regions. *: estimates are significant at 1%

on the drug price index and donut hole price index, 0.78. The point to notice is that the across year correlation in LIS eligibility is the smallest of all plan characteristics, 0.6. This lower persistence is suggestive that the strategic aspects of bidding matter for LIS eligibility. Note that a pure cycling story would predict a negative correlation, and we see a positive correlation. This positive sign could reflect across year correlation in cost. Low cost plans today are likely to be low cost plans tomorrow. The evidence on cycling comes from the comparison in correlation with other plan characteristics that show much higher persistence.

5 Risk Selection

The proposed supply side model abstracts away from the issue of risk selection. Medicare uses a risk adjustment mechanism designed to thwart adverse selection. Payments to plans are adjusted up or down based on the risk scores of a plan's enrollees. The risk score $r_i > 0$ for enrollee i is calculated using Medicare's risk adjustment formulary, which predicts drug expenditures based on demographic (age, gender) and disease factors (kidney failure, diabetes, etc). The risk score for the average expenditure enrollee is normalized to $r_i = 1$. Plan j's risk score r_j is the average risk score of its enrollees. Under the risk adjustment mechanism, the payment to the plan is $b_j r_j$. For example, the payment is 10% above the bid for a plan with a risk score of 1.1.

With a perfect risk adjustment mechanism, plans price with respect to the cost of the average Medicare Beneficiary; it does not matter whether they attract a low risk or high risk pool of enrollees. However, with an imperfect risk adjustment mechanism it is possible for plans to cream skim the relatively low cost enrollees. There is evidence of cream skimming occurring in Medicare Advantage markets (Brown et al., 2011). Presumably it also happens in Medicare Part D. Plans cream skim by tweaking drug formularies to target specific disease categories (Hsu et al., 2009; McAdams and Schwarz, 2007). For example, the plan can favorable cover diabetic drugs that appeal to the relatively healthy diabetics and place restrictions on drugs desired by the higher cost types. If a plan is successful in cream skimming it will have a pool of enrollees with cost lower than that predicted by the risk adjustment formula.

This idea can be formalized by defining a *selection factor* a_i to measure the deviance of an enrollee's expected cost from that predicted by the risk adjustment formula. Denote the

year	All plans	Stand-alone Part D	Stand-alone per person per month
2006	-\$2,588mil	-\$1,228mil	-\$5.53
2007	-\$599mil	-\$204mil	-\$0.95
2008	-\$78mil	\$100mil	\$0.51
2009	-\$795mil	-\$500mil	-\$2.01
2010	-\$900mil	-\$395mil	-\$1.63

cost to plan j of enrolling a person with risk score r_i and selection factor a_i as,

$$c_{ij} = c_j(r_i + a_i)$$

The term c_j is a plan specific cost scalar. With $a_i > 0$ the plan selects relatively high cost enrollee, and $a_i < 0$, relatively low cost. If $a_i = 0$, the enrollee cost exactly matches the cost predicted by the risk adjustment formula. Averaging r_i and a_i for plan j the cost per enrollee is $c_j(r_j + a_j)$, where r_j and a_j are the average risk score and selection factor of the enrollees.

Besides risk adjustment payments, Medicare also uses a profit/loss sharing mechanism called risk corridors. If a plan attracts enrollees with cost greater than that predicted by the risk adjustment formula ($c_{ij} > c_j r_i$), medicare makes a payment to the plan to compensate for the loss of profits. Likewise, if costs are lower ($c_{ij} < c_j r_i$), Medicare deducts payments to claw back profits. The fraction of sharing is denoted by the parameter $\theta \in [0, 1)$. In early years profit sharing was high $\theta = 0.8$, but was reduced in later years to $\theta = 0.5$. The marginal rate of sharing depends on the degree of losses/profits. For small losses/profits there is no sharing, then increases at a profit/loss threshold of +/- 5%, and again at +/- 10%. The reconciliation process for risk sharing takes 2 years.

With linear risk corridors, the profit function for a plan with bid b_j, risk score, r_j and selection factor a_j is

$$\pi_{jmt} \quad = \quad [(b_j r_j + \theta c_j a_j) - c_j(r_j + a_j)]\, s_{jmt}(\mathbf{b}) M_{mt} \tag{29}$$

CMS recently released data on risk sharing reconciliation, aggregated at the insurer level. It reports the money clawed back or paid out by the risk corridors. In the notation of this model, $\sum_{j \in f} \theta c_j a_j$ for all plans offered by insurer f. Aggregated across all plans, CMS clawed back substantial amounts of money which suggests plans were quite successful in cream skimming low a_j types. The following table reports statistics. Negative signs indicate payments made from plans to CMS, or favorable risk selection. In the first year, there were very large risk corridor payments. For stand-alone part d plans, the amount is over $5 per person per month, given the average bid is around $90, that is a large amount. It reduced in later years, and actually flipped signs in 2008, indicating plans suffered from adverse selection. That was the year CMS introduced an updated risk adjustment formula. But after that, in 2009 and 2010, plans achieved favorable selection again.

5.1 Does Random Assignment mitigate cream skimming?

The above evidence suggests plans are quite successful cream skimming favorable risk. Key to cream skimming is that enrollees make active choices to sort into plans. For randomly assigned LIS enrollees, plans may not be able to cream skim the favorable risk, rather they will get the average risk. I test this hypothesis using insurer level data on risk corridor payments. Let $\theta_t c_{ft} a_{ft}$ denote the measure of risk corridor payments for insurer f in year t in per enrollee per month terms. I specify the following diff-in-diff regression.

$$\theta_t c_{ft} a_{ft} - \theta_{t-1} c_{ft-1} a_{ft-1} = \beta(LISfrac_{ft} - LISfrac_{ft-1}) + \alpha_f + \alpha_t + \epsilon_{ft}$$

The regressor $LISfrac_{ft}$ is the fraction of insurer f's enrollees in an LIS eligible plan. α_f and α_t are insurer and time fixed effects. The diff-in-diff uses across time variation in whether in an insurer is moving in or out of LIS eligibility to see how its selection changes in response to LIS eligibility. The hypothesis is that an increase in LIS enrollees increases the selection factor a_f, due to the fact that cream skimming practices will be less effective. The coefficient β is very large and positive: 8.77 with standard error 2.25. In dollar terms, this means in literal terms that an insurer going from no LIS eligible plans to all plans being LIS eligible, will experience an $8.77 difference in risk corridor payments per enrollee per month. Using $\theta = 0.5$ as an approximate measure[4] the coefficient implies a $17.54 dollar increase in adverse selection cost $c_{ft} a_{ft}$ when an insurer shifts from no LIS plans to all LIS plans. It is rare for an insurer to go from no-LIS to all LIS, however it is quite common to go from all LIS to no-LIS (over 25% of all insurers), which translates into a $17.54 reduction in cost due to its ability to favorably select. This is a large figure given the average bid is about $90.

The next part of the project is to tie this back into to the supply-side model as a discontinuity in cost curves when crossing into LIS eligibility. The FOC of the profit equation with selection in equation 29 will be different because a_j terms distort pricing.

[4] I can't figure exactly because risk corridors are phased-in.

6 Data

We collected publicly available data from the Center for Medicare and Medicaid Services (CMS) on plan level enrollment and bids for all stand alone part D plans offered since the the programs inception in 2006 through 2009. We also purchased detailed data on plan characteristics from CMS. There are four files. The formulary file lists all drugs on a plan's formulary, the beneficiary cost file describes cost sharing rules, the pharmacy network file lists all preferred and non-preferred pharmacies, and the pricing file reports average drug transaction prices for every drug and plan. The pricing file was first published in 2009, the other files are available in all years. Specifically, prices are the average transaction price net of all rebates for a 30 day supply filed at the plan's preferred pharmacies in Q3 2009.[5] They are used to calculate enrollee drug expenditures. It's worth noting that enrollees may not know the exact drug price during the enrollment period because drug prices and rebates vary throughout the year, and prices reported by the "Plan Finder" tool on Medicare's website are not necessarily accurate.[6]

6.1 Enrollment and Premiums

Across the four years and 39 regions, 75 insurers offer stand alone Part D plans. The market penetration of insurers is quite skewed; 18 national firms offer plans in all states,[7] while 44 regional insurers offer plans in just one market. Insurers offer an average of 2.5 plans in a market, with most offering 1, 2, or 3 plans.[8] At least one must be a basic plan.

Enrollees have lots of choices. A typical enrollee can choose from over 40 plans offered by about 20 insurers. Table 4 shows the average number of insurers and plan offerings in a

[5] Plan's report all transactions, called Prescription Drug Events (PDE) to CMS. A PDE includes information on prices and all rebates/discounts with manufacturers, wholesalers, and pharmacies. Rebate information is proprietary, only the net price is available to researchers. Some pharmacies charge a dispensing fee, on the order of a couple dollars, that may or may not be included in the net drug price. Our data are based on PDE records.

[6] Every two weeks plans are required to submit a separate pricing file to CMS that is used in the plan finder database. The database is not constructed from PDE records. If the plan never submits a price, the finder reports a price 30% below the average wholesale price for generic drugs, 10% below for brand name drugs. Even if a plans submits a price, it may not get updated every 2 weeks, so the prices can be outdated. Recently, Medicare began reporting survey results that ask enrollees to rate the accuracy of drug prices paid at the pharmacy compared to price reported on the plan finder. Many plans get very poor ratings.

[7] excluding 5 markets for US territories with very small penetration

[8] There are a few exceptions where insurers offer more than 3 plans, up to a maximum of 9. The most notable is in 2006 when United Healthcare, the market leader, offered 5 plans. Two of these plans were sponsored by Pacificare, which was acquired by United Healthcare. United Healthcare consolidated these 5 plans into 3 in 2007.

market. The entries are broken down by year and plan segment: Enhanced and Basic.

Despite the large number of insurers, the Part D market is highly concentrated. Table 5 reports various measures of firm concentration (1-firm concentration, $c1$, 2-firm $c2$, 4-firm $c4$, and Herfindahl-Hirschman index) averaged across markets. The top firm (United Healthcare for most markets) commands an average 36% market share, while the top 4 firms, 73%. The HHI averages 2376, which according to the Department of Justice guidelines falls into the "moderately concentrated" category.[9] Overtime, the markets have become less concentrated but still fall into moderately concentrated levels.

Table 6 reports national enrollment as a percentage of all eligible Medicare beneficiaries ($\approx$ 42 million) and monthly premiums. The table divides shares into three categories: stand alone part D plans, bundled MA+Part D, and stand alone Medicare Advantage plans. Stand alone Part D enrollment has been stagnant since program inception in 2006, while monthly premiums have risen dramatically; average premiums rose about 30% between 2006 and 2009. Meanwhile, both stand alone MA and bundled MA+PartD, plans have experienced an increase in enrollment. Tables 7 reports enrollment and premium statistics separate by basic and enhanced plans. Basic plans attract about 3 times as many enrollees and charge 30% lower premiums as compared to enhanced plans.

Table 8 reports more detailed summary statistics for both basic and enhanced plans at the market level. Note the large variation in premiums for both plan segments as well as the variation and skew in enrollment figures. The average basic plans enrolls 1% of Medicare beneficiaries, while the largest, upwards of 18%. The average enhanced plan has a smaller share, .4%, but the distribution also exhibits a large skew.

Every year Medicare announces the official average bid amount $\bar{b}_t^{basic}$ and average premium $\bar{b}_t^{basic} - \lambda_t \bar{b}_t^{basic}$. These figures, along with the calculated subsidy fraction λ_t, are reported in table 9.[10] Table 10 reports summary statistics on plan bids.

[9] A market with a HHI between 1800 and 2500 is considered moderately concentrated.

[10] The official average bid and average premium are reported by Medicare. We also observe the basic and enhanced component of the bid for all stand alone Part D plans. We do not have complete data on MA+Part D bids due to further complications in the rules for subsidizing MA+Part D plans. Instead, we use enrollment figures from MA+Part D plans and the subsidy formula (equation 10) to calculate an average basic bid for MA+Part D plans. Since MA+Part D plans are included in the outside option, the average bid is a sufficient statistic for us to properly calculate marginal cost and perform counterfactuals. Because of the phase-in of the weighted average bid method, we are missing data that would allow us to calculate marginal cost in 2006 through 2008.

6.2 Plan Characteristics

Our primary plan characteristic variables measure the generosity of plans coverage. Our first variable is the deductible. The second and third are intended to measure the generosity of coverage in the initial coverage and donut hole regions. The challenge is taking our rich drug-level data and converting it into a meaningful plan-level characteristic.

We construct price indices for the top 100 most popular drugs ranked by prescriptions filled.[11] Our first price index reflects the out of pocket cost for an enrollee to fill a 30 day supply for a basket of the 100 drugs when they are in the initial coverage zone. Our second price index reflects out of pocket costs in the donut hole. The basket of drugs evenly weights each drug (1/100th). While there may be drug-by-drug idiosyncratic variation in a plan's out of pocket prescription drug cost, this measure captures a plan's average cost across drugs. Constructing out of pocket costs is straightforward for drugs covered by a copay. For drugs covered by coinsurance, it is necessary to know the price of the drug. We use the 2009 pricing file. For off-formulary drugs, enrollees do not receive coverage, therefore the out of pocket cost is the full retail price. We set the retail price to the average price in the region.[12] We do not construct a price index for the catastrophic region because there is virtually no variation across plans.[13]

There are three sources of variation in the price indices: copay and coinsurance rates, negotiated drug prices, and formulary composition. Table 12 reports statistics on out of pocket price indices for the top 100 drugs and separate indices for brand and generic. Most of the variation in the donut hole is between enhanced plans that fill the donut hole and basic plans with no coverage. There is more variation in the initial coverage zone than in the donut hole. The source of this variation comes from differences in copay rates. Comparing brand and generic drugs, the variation is larger for brand name drugs. Its also interesting to note trends across time. Average donut hole prices remain steady, while out of pocket prices

[11] CMS published a report ranking the top 100 drugs by number of prescriptions filled by Part D enrollees in 2006. Rankings by cost are quite different. For example, the generic drug FUROSEMIDE is number 1 by prescriptions filled and 98 by cost.

[12] Since the base price includes rebates and discounts, we are probably understating retail prices by using negotiated prices. We use the average national price in rare cases where a region price does not exists. For the years 2006 to 2008 we construct the price indices in the same manner using 2009 prices. For plans that did not exists in 2009, we use average regional prices. Drug prices, coinsurance rates, and copays differ across preferred, non-preferred, and mail order pharmacies. All of our calculations are based on preferred pharmacies.

[13] The Part D regulations do not allow plans to use a tiered copay/coinsurance structure in the catastrophic region. Out of pocket payments are capped at $5 per prescription or 5% of drug cost. There is little variation across plans. Moreover, few enrollees, only 8% in 2006 reached the final tier, and of that group, they are over-represented by the low income subsidy enrollees who pay zero in the catastrophic region.

in the initial coverage zone fall across years.

Figure 6 shows a histogram of negotiated prices for all drugs in 2009. To compare across drugs, we record prices as percent deviations from the drug's average price. Notice there is a lot of price dispersion; it contributes to the variation in our price indices. To give a sense of magnitude, 10% of drugs are priced 25% below the average, and 10% are priced 15% above average. The dispersion is quite remarkable considering these are perfectly homogenous products.

We measure formulary comprehensiveness by counting the number of top 100 drugs included on a formulary. We also break this list down by brand and generic medication; there are 42 brand name medications and 58 generic. Table 11 reports statistics on formulary comprehensive. On average plans cover most of the drugs (more than 90%), but there is significant variation that appears to have grown over the years, indicating the plans are more differentiated now than in 2006. There is little difference between enhanced and basic plans. Insurers typically share formularies for their plans. Across all four years and regions there are 6679 plans and 400 formularies. Th only intended to illustrate a source of variation in the price index; we do not includes these drug counts as a separate plan characteristic.

Across the universe of Part D drugs, over 5400, there is a lot of idiosyncratic variation in formularies. Figure 5 depicts a snapshot of formularies in 2009. Gaps in the formularies show that less comprehensive formularies are not strict subsets of more comprehensive formularies. With so many non-overlapping formularies, each enrollee is likely to find a plan tailored his individual drug regimen. This suggests enrollees sacrifice very little in terms of choice when plans use formulary restrictions as a bargaining chip with drug manufacturers.

From our data on pharmacy networks, we construct a measure of network coverage by counting the number of in-region network pharmacies per eligible Medicare beneficiary in the region. We group preferred and non-preferred pharmacies because many plans do not make a distinction.

6.3 Basic vs Enhanced Plan Characteristics

Figures 7, 8, and 9 display side-by-side histograms of coverage characteristics for basic and enhanced plans in 2009. They illustrate the relative location in characteristic space for the two categories. The differences are relevant because they drive the substitution patterns in our counterfactual that introduces a basic government option. The plans are highly differentiated with respect to the monthly deductible. Most basic plans have the maximum deductible ($295 in 2009), while most enhanced plans have a $0 deductible. No systematic difference emerges in the histograms for the price index in the initial coverage zone, but there are differences in the donut hole price index. Almost no basic plans have a price index below 100, while about 50% of the enhanced plans fall below the $100 thresholds. The price index is lower for these enhanced plans because they provide some coverage in the donut hole.

6.4 Heterogeneity in Preferences: Random Coefficients

Our most flexible demand specification includes random coefficients on the monthly premium, deductible, and out of pocket price indices. As in a typical demand system, the random coefficient α_i captures heterogeneity in consumers' marginal utility over income, driven by differences in income and price sensitivity. Two other factors affect the distribution of α_i: the low income subsidy and late enrollment penalty. The low income subsidy truncates α_i. For example, it is zero for those receiving a 100% subsidy and half of what it would otherwise be for an individual receiving a 50% subsidy. For an enrollee subject to the late enrollment, penalty—1% per month not enrolled— α_i increases in proportion to the duration of late enrollment. Combined, the low income subsidy and late enrollment penalty increase the variance on the distribution of α_i. They may also skew the distribution, but we lack data on these populations to estimate higher order moments of the distribution.

Heterogeneity in preferences for the deductible and out of pocket price indices are driven by heterogeneity in enrollee's health status and risk aversion. All enrollees (weakly) prefer a lower deductible, but "Healthy" enrollees have a relatively low (in magnitude) marginal utility with respect to the deductible because their drug expenditures are unlikely to exceed the deductible. "Sick" enrollees have a higher marginal utility because they would expect to spend through the deductible with certainty. By the same reasoning, health status affects preferences over out of pocket prices in the initial coverage and donut hole regions. It's worth noting, enrollees care not only about the cost sharing rates in the marginal tier of the tariff schedule they fall into (deductible, initial coverage, donut hole, or catastrophic), but also tiers they surpassed. Marginal utility only diminishes for the higher tiers that they are unlikely to enter. More risk averse enrollees place a higher preference on the deductible and price indices. Like the premium, the low income subsidy truncates the distributions towards zero.

We use a parsimonious Normal distribution over the random coefficients with a block diagonal covariance matrix. However, this distribution may not be appropriate for three reasons. First, in a model of demand for drugs with a kinked tariff schedule, the distribution of drug expenditures will exhibit bunching and gaps around the kinks (Marsh, 2010). In our model of plan demand, this translate into gaps and mass points in the distribution of random coefficients. Second, we would expect correlation in the random coefficients across the tiers of the tariff schedule because they all relate to expenditures of money. For a risk neutral enrollee, a dollar spent on the premium is worth exactly the same as an (expected) dollar spent in the later tiers. This implies an enrollee's marginal utility over characteristics monotonically decreases moving down the tariff schedule. The premium should have the highest coefficient because it is paid with certainty, and the donut hole should have the lowest coefficient since it is not necessarily reached with certainty. Third, if consumers are risk averse, a risk premium is built into their preferences which further complicates the correlation structure of random coefficients. The distribution over the marginal tier that an enrollee enters matters for the risk premium. For example, a consumer that exceeds the deductible with certainty places a risk premium on the coefficient in later tiers and no risk premium on the deductible coefficient. For estimation, we experimented with specifications that accommodate these three qualifications (correlation in random coefficients, discrete type models (Berry and Jia, 2010), and mixtures models with Normal and discrete distributions) but could not obtain sensible results. In principle these richer models may be identified with aggregate data (**?**), but accessing consumer level data, as in Abaluck and Gruber (2011), would help with identification. We leave this as future work. Nonetheless, the focus of this paper is on the the supply-side for which the demand model is a means to obtain reasonable price elasticity estimates. With our more parsimonious random coefficient model, we estimate sensible elasticities and ameliorate the standard criticisms of non-random coefficient discrete choice models.

6.5 Heterogeneity in Preferences: Idiosyncratic Preferences

The idiosyncratic logit error terms, ϵ_{ij}, reflect unobserved heterogeneity in preferences that are not otherwise captured by random coefficients. There are several reasons we believe they should enter our demand specification. Drug purchasing patterns are likely the primary source of idiosyncratic preferences. Enrollees have stronger preferences for plans that cover their specific drug regimen at low out of pocket prices. Thus, drug-by-drug idiosyncratic differences in formulary composition and copay/coinsurance rates generate idiosyncratic preferences. Figure 5 illustrates idiosyncratic differences in formulary composition; a similar visual representation of copay/coinsurance rates shows the same pattern. Marketing activities and pharmacy networks may also contribute to idiosyncratic preferences. Using the examples from before, AARP members might have stronger preferences for AARP endorsed plans, and Walmart customers might prefer plans that contract with Walmart. It is also worth noting low income subsidy enrollees who accept random assignment to plans. Their behavior can be rationalized in our model by attributing the random assignment to draws from the distribution of ϵ_{ij}.

Ackerberg and Rysman (2005) and Berry and Pakes (2007) have critiqued this model because the dimensionality of ϵ_{ij} increases as products are added. The added term ensures consumers benefit from the introduction of a public option, even if it has a high premium, undesirable average characteristics, and there is no competitive response by existing plans. We justify the extra ϵ_{ij} term for the government plan by assuming its formulary and co-pay/coinsurance rates will exhibit the same sort of idiosyncratic differences found amongst the privately offered plans. These difference are a likely outcome because the legislation permits the government to bargain with drug manufacturers using restrictive formularies and tiered copays, just like private plans.

7 Estimation Results

The following table presents demand estimates for the LIS and non-LIS segments of the market. They are estimated separately on the two segments of the market in a flexible manner with no restrictions placed across segments.

Demand estimates

7.1 PRELIMINARY MC estimates

The following tables report enrollment weighted averages for marginal cost and markups for 1 national firm in the sample that offers 102 plans, 9 of which are priced at or below the LIS threshold. The results are based on the assumption that profits are proportionally allocated to the basic and enhanced component.

8 Counterfactual

Removing LIS population from the market. Reweighting LIS threshold.

[TO BE CONDUCTED]

	LIS population Logit	Non LIS population Logit	LIS population RC	Non LIS population RC
Premium (α)	-.059	-.065	-.213	-.106
	(.005)	(.007)	(.045)	(.034)
Std Dev(premium)			.089	.075
			(.017)	(.012)
LIS Subsidy ($\kappa\alpha$)	.054		.193	
	(.024)		(.034)	
Std Dev(LIS Subsidy)			.105	
			(.040)	
deductible/12	-.044	-.067	-.051	-.372
	(.009)	(.005)	(.012)	(.061)
Std Dev(deductible)			.041	.207
Initial coverage			(.451)	(.169)
Price index	-.044	-.031	-.039	-.028
	(.005)	(.004)	(.006)	(.008)
Donut Hole				
Price index		-.081		-.046
		(.016)		(.035)
Std Dev(index)				.039
				(.030)
pharm per eligible (x1000)	1.55	1.85	1.88	3.66
	(0.29)	(0.22)	(0.45)	(0.63)
LIS eligible plan (indicator)	4.27		3.26	
	(0.53)		(0.69)	
N obs	1575	1575	1575	1575
N sims	—	—	400	400
Gmm Obj Func	454.31	428.84	310.24	254.97

Standard errors in parentheses. Correlation amongst RC=0.985

	no restriction	Restriction
Enhanced	80.03	79.98
Basic Non-LIS	82.10	81.79
LIS	53.85	75.13
all Basic	60.72	76.74

Lerner Markup % b-mc/b

	no restriction	Restriction
Enhanced	14.21	14.25
Basic Non-LIS	5.93	6.28
LIS	35.43	9.97
all Basic	28.26	9.07

9 Conclusion

[TO BE WRITTEN]

References

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

D. Ackerberg and M. Rysman. Unobserved Product Differentiation in Discrete-choice Models: Estimating Price Elasticities and Welfare Effects. *RAND Journal of Economics*, 36 (4):771–788, 2005.

S. Berry. Estimating Discrete-choice Models of Product Differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994.

S. Berry and P. Jia. Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics*, 2(3):1–43, 2010. ISSN 1945-7669.

S. Berry and A. Pakes. The Pure Characteristics Demand Model. *International Economic Review*, 48(4):1193–1225, 2007.

S. Berry, J. Levinsohn, and A. Pakes. Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890, 1995.

J. Brown, M. Duggan, I. Kuziemko, and W. Woolston. How Does Risk Selection Respond to Risk Adjustment? Evidence from the Medicare Advantage Program. *NBER Working Paper*, 2011.

B Handel. Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review*, forthcoming, 2013.

F. Heiss, D. McFadden, and J. Winter. Who failed to enroll in Medicare Part D, and why? Early results. *Health Affairs*, 25(5):w344, 2006.

F. Heiss, D. McFadden, and J. Winter. Mind the gap! Consumer Perceptions and Choices of Medicare Part D Prescription Drug Plans. *Research Findings in the Economics of Aging*, pages 413–481, 2010.

J. Hsu, J. Huang, V. Fung, M. Price, R. Brand, R. Hui, B. Fireman, W. Dow, J. Bertko, and J. Newhouse. Distributing $800 Billion: An Early Assessment of Medicare Part D Risk Adjustment. *Health Affairs*, 28(1):215–225, 2009.

J. Ketcham, C. Lucarelli, E. Miravete, and M. Roebuck. Sinking, Swimming, or Learning to Swim in Medicare Part D. *Working Paper*, 2011.

J.R. Kling, S. Mullainathan, E. Shafir, L. Vermeulen, and M. Wrobel. Comparison Friction: Experimental Evidence From Medicare Drug Plans. *Quarterly Journal of Economics*, forthcoming, 2012.

C. Lucarelli, J. Prince, and K. Simon. The Welfare Impact of Reducing Choice in Medicare Part D: A Comparison of Two Regulation Strategies. *International Economic Review forthcoming*, 2011.

J. Lustig. Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare. *Working Paper Boston University*, 2010.

C. Marsh. Estimating Health Expenditure Elasticities Using Nonlinear Reimbursement. *Working Paper University of Georgia*, 2010.

D. McAdams and M. Schwarz. Perverse Incentives in the Medicare Prescription Drug Benefit. *Inquiry*, 44(2), 2007.

D. Miller and J. Yeo. The Consequences of a Public Health Insurance Option: Evidence from Medicare Part D Prescription Drug Markets. *Working Paper Clemson University, Singapore Management University*, 2012a.

D. Miller and J. Yeo. Estimating Dynamic Discrete Models of Product Differentiation: An Application to Medicare Part D with Switching Costs. *Working Paper Clemson University, Singapore Management University*, 2012b.

K. Nosal. Estimating Switching Costs for Medicare Advantage Plans. *Working Paper Mannheim*, 2012.

R. Town and S. Liu. The Welfare Impact of Medicare HMOs. *RAND Journal of Economics*, pages 719–736, 2003.

A Data Tables and Figures

Table 3: **Part D Basic Benefit Structure**

Tariff Level	Name	MC of Drugs	Expenditure Range
1	Premium	—	0
2	Deductible	1	[0,250]
3	Initial Coverage	0.25	[250,2250]
4	Donut Hole	1	[2250,5100]
5	Catastrophic	0.05	> 5100

The marginal cost of drugs represents the marginal out of pocket payment an enrollees pays for each dollar of drug expenditures. These are the 2006 thresholds. They increase each year to keep the program on budget.

Table 4: **Number of Insurers and Plan Offerings per Market**

Year	All Plans	Basic	Enhanced	Insurers
2006	37	21	21	14
2007	46	23	23	20
2008	48	24	24	18
2009	43	19	19	18
All	43	22	22	17

Number of plan and insurer offerings averaged across markets. There are a total of 39 markets.

Table 5: **Firm Concentration Measures by Market**

year	$c1$	$c2$	$c4$	Hirf
2006	37.5	58.2	75.7	2587
2007	37.7	57.6	75.3	2558
2008	34.0	52.0	70.0	2154
2009	34.6	50.1	70.2	2196
All	36.0	54.7	72.9	2376

Firm concentration measures averaged across markets. $c1$, $c2$, $c4$ represent the inside market share of the top 1,2 and 4 insurers in a market. According to DOJ guidelines, a HHI index between 1800-2500 is “moderately concentrated”. Above “2500” highly concentrated.

Table 6: **National Enrollment and Premiums**

	National Enrollment (%)			Part D Avg. Monthly Premium ($)	
Year	Part D	MA+Part D	MA	enrollment weighted	Unweighted
2006	46.3	19.4	2.5	26.67	37.36
2007	43.7	19.8	3.4	27.40	36.68
2008	41.6	22.7	4.1	29.95	39.86
2009	46.1	28.2	4.3	34.71	45.27

The enrollment percentage is the ratio of enrollment counts to the total number of eligible Medicare Beneficiaries in the nation (growing from about 40 million in 2006 to 45 million in 2009). Part D and MA+Part D are mutually exclusive groups, but MA enrollees may also enroll in a stand-alone Part D plan. The monthly premium figures are reported as enrollment weighted and nonweighted averages.

Table 7: **National Enrollment and Premiums: Basic vs Enhanced Plans**

	Part D National Enrollment (%)		Part D Avg. Monthly Premium ($)			
			Enhanced		Basic	
Year	Enhanced	Basic	Enrollment weighted	Unweighted	Enrollment weighted	Unweighted
2006	9.9	36.4	34.98	41.62	24.40	34.27
2007	9.0	34.8	40.37	45.51	24.05	28.69
2008	9.4	32.2	40.93	49.49	26.77	30.01
2009	11.9	34.1	45.86	55.72	30.82	33.51

The enrollment percentage is the ratio of enrollment counts to the total number of eligible Medicare Beneficiaries. The monthly premium figures are reported as enrollment weighted and nonweighted averages.

Table 8: **Enrollment and Premium Summary Statistics**

	Monthly Premiums ($)									
	Basic Plans					Enhanced Plans				
	mean	s.d.	min	max	obs	mean	s.d.	min	max	obs
2006	33.9	11.9	1.9	70.8	825	41.6	12.3	4.9	99.9	606
2007	27.6	6.4	1.9	49.0	914	45.4	16.4	17.1	135.7	890
2008	29.4	9.4	2.6	72.0	897	49.5	22.2	12.9	107.5	929
2009	33.7	10.0	1.0	112.7	734	55.6	22.0	10.3	136.8	884
	Region Level Market Shares (%)									
	Basic Plans					Enhanced Plans				
	mean	s.d.	min	max		mean	s.d.	min	max	
2006	1.3	2.2	0	18.0		0.5	0.9	0	10.2	
2007	1.1	2.1	0	18.1		0.3	0.7	0	6.6	
2008	1.1	1.9	0	18.5		0.3	0.9	0	10.8	
2009	1.1	1.9	0	17.1		0.4	0.7	0	8.6	

The top panel reports summary statistics on monthly premiums for basic and enhanced plans. The bottom panel reports market share statistics, expressed as a ratio of the total number of eligible Medicare Beneficiaries in the plan's region. These are not inside market shares.

Table 9: **Official Average Bid and Premium**

year	avg premium	avg bid	λ
2006	32.20	92.30	0.651
2007	27.35	80.43	0.660
2008	27.93	80.52	0.653
2009	30.36	84.33	0.640

The average bid and avg premium values are collected from official Medicare data releases. They are not calculated from our bid and enrollment data. For 2009 Medicare used equation 10 to calculate the average bid. It is based on the basic component of bids and includes Part D and select MA+Part D plans. We cannot replicate Medicare's calculation because we lack data on the basic component of bids for MA+Part D plans. In 2006, the bids of stand-alone Part D plans and new MA+Part D plans were evenly weighted and were based on lagged enrollment for MA plans that existed in 2005 and became MA+Part D plans. The shift to lagged enrollment weighting was phased in for 2007 and 2008. We calculate the subsidy ratio as $\lambda = \frac{avgbid - avgpremium}{avgbid}$.

Table 10: **Bid Summary Statistics**

	Monthly Bid			
	mean	s.d.	min	max
2006	97.3	12.6	62.0	160.0
2007	89.5	15.3	55.0	188.8
2008	92.2	19.9	55.2	160.1
2009	99.6	20.7	55.0	190.8
	Monthly Basic Bid Component			
	mean	s.d.	min	max
2006	92.5	11.2	62.0	127.3
2007	81.5	7.5	55.0	111.4
2008	83.9	12.5	55.1	133.8
2009	90.7	15.0	55.0	166.7
	Monthly Enhanced Bid Component			
	mean	s.d.	min	max
2006	4.8	9.1	0	55.3
2007	8.0	11.8	0	96.4
2008	8.3	11.8	0	55.0
2009	8.9	10.7	0	55.8
	Enhanced/Basic Bid Ratio			
	mean	s.d.	min	max
2006	.056	.118	0	.716
2007	.097	.140	0	1.32
2008	.094	.134	0	.670
2009	.096	.117	0	.764

This table reports summary statistics on bids b, the basic component of bids b^{basic}, and the enhanced component of bids $b^{enhanced}$, where $b = b^{basic} + b^{enhanced}$. The bottom panel reports the ratio of the enhanced component to the basic component, corresponding the parameter γ in our model. This table only includes the bids of stand-alone Part D plans.

Table 11: **Formulary Comprehensiveness**

	num	top100		Brand(top42)		Generic(top58)	
	plans	avg	s.d.	avg	s.d.	avg	s.d.
2006	1446	90.5	5.9	38.7	3.9	51.5	3.2
2007	1909	92.1	6.0	39.0	3.4	53.2	3.1
2008	1877	89.2	7.5	37.6	4.1	51.6	3.9
2009	1650	86.8	9.1	35.6	5.5	51.2	4.2

This table reports statistics on the number of top 100 drugs on a formulary. The top 100 drugs are ranked by prescriptions filled. The table also breaks the statistics down for the 42 brand name drugs and 58 generic drugs composed in the top 100.

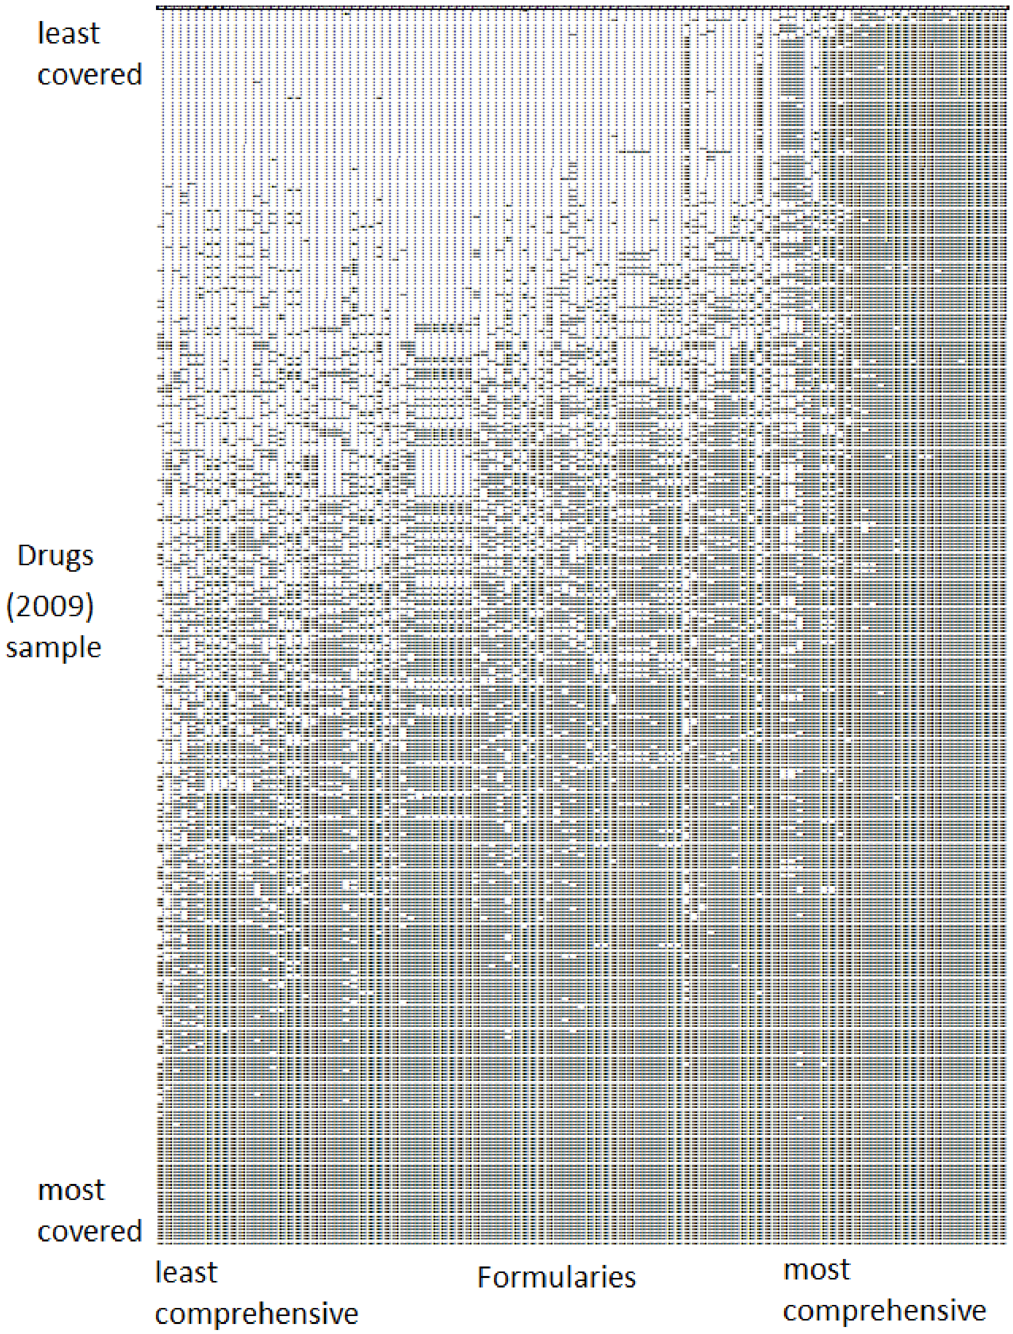

Figure 5: **2009: Formulary Snapshot**
The figure depicts the drugs covered on a formulary for a 1/15 sample of the 5400 drugs and a 1/3 sample of the 315 insurers offering plan's in 2009. This includes both Part D and MA+Part D insurers. Whitespace indicates off-formulary drugs. Notice the idiosyncratic differences in formularies; that is, less comprehensive formularies are not strict subsets of more comprehensive formularies. Generic drugs, with multiple manufacturers, are counted once. Different packages sizes of the same drug appear as distinct drugs.

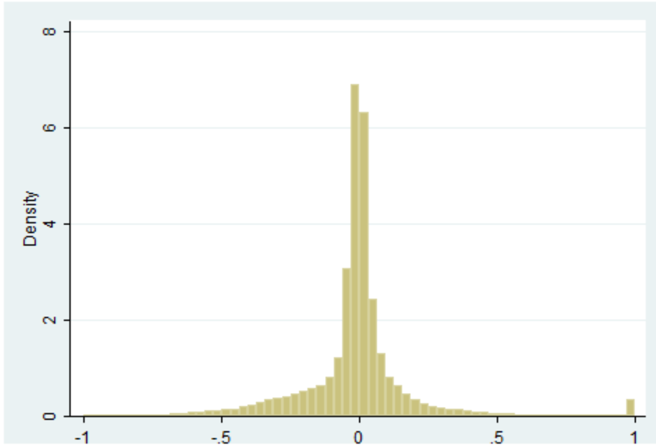

Figure 6: **2009 Drug Price Dispersion, Deviation from Avg Negotiated Drug Price**

This histogram displays the price dispersion in negotiated drug prices. The x-axis is the percent difference in price p_{jd} for a drug d offered by plan j from the average price for that same drug $\bar{p}_d$ ($\frac{p_{dj}-\bar{p}_d}{\bar{p}_d}$). There are 16,781,151 drug-plan observations, 5330 drugs defined by a unique NDC (national drug code), and 4228 plans (including MA+Part D plans).

Table 12: **Out of Pocket Drug Price Indices**

	Initial Coverage Zone								
	num	top100		Brand(top42)		Generic(top58)		top100	
	plans	avg	s.d.	avg	s.d.	avg	s.d.	min	max
2006	1446	67.9	9.8	70.6	22.1	52.9	10.7	41.0	90.7
2007	1909	62.7	9.9	48.7	15.7	56.7	18.0	44.1	83.5
2008	1877	58.6	10.9	56.2	16.3	60.4	16.0	29.0	78.7
2009	1650	55.6	11.0	71.9	18.4	50.5	9.4	24.0	73.0
2010	-	-	-	-	-	-	-	-	-

	Donut Hole								
	num	top100		Brand(top42)		Generic(top58)		top100	
	plans	avg	s.d.	avg	s.d.	avg	s.d.	min	max
2006	1446	101.2	5.9	185.3	17.4	68.6	10.7	54.9	107.6
2007	1909	99.8	6.0	188.1	16.5	86.9	10.2	53.0	113.1
2008	1877	100.0	5.2	167.6	44.3	84.0	16.3	78.9	113.1
2009	1650	99.9	5.8	160.4	39.8	64.1	23.9	82.0	113.1
2010	-	-	-	-	-	-	-	-	-

The top panel reports summarary statistics for the out of pocket drug price indices for top 100 drugs in the initial coverage zone. The bottom panel reports for the donut hole. The price is what a consumer pays out of pocket for a 30 day supply. For on-formulary drugs, we first locate the copay or coinsurance rate corresponding to that drug from the beneficiary cost file. The out of pocket price is either the copay or coinsurance rate times the negotiated drug price found in the 2009 Q3 pricing file. For 2006-2008 we match to the plan's negotiated price in 2009. If a plan did not exist in 2009, we use the average negotiated price in the region, or if there aren't enough observations, the national average price. For off-formulary drugs, the consumer pays full retail price which we set equal to the regional or national average negotiated price in 2009. All drugs are evenly weighted in the index.

Table 13: **Network Pharmacies**

	# Network Pharmacies per Eligible Beneficiary			
	mean	s.d.	min	max
2006	.00141	.00024	0	.00376
2007	.00137	.00021	0	.00193
2008	.00135	.00020	0	.00189
2009	.00136	.00018	0	.00186

Summary statistics about the number of network pharmacies per eligible Medicare Beneficiary. We include both preferred and non-preferred pharmacies because many plans don't make a distinction. These are brick-and-mortar pharmacies located in the region. We exclude out of region network pharmacies and mail order pharmacies.

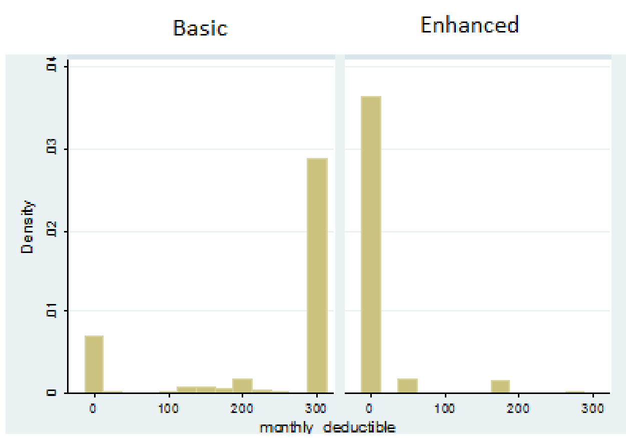

Figure 7: **2009: Monthly Deductible Histogram**
This histogram compares the deductible of basic and enhanced plans. Notice most basic plans have a $295 deductible and most enhanced plans have a $0 deductible.

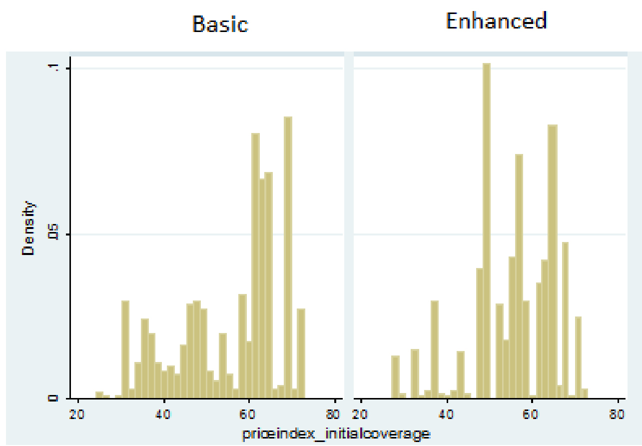

Figure 8: **2009: Initial Coverage Price Index Histogram**
This histogram compares the initial coverage zone out of pocket price index of basic and enhanced plans. Notice there is no discernible difference between basic and enhanced plans.

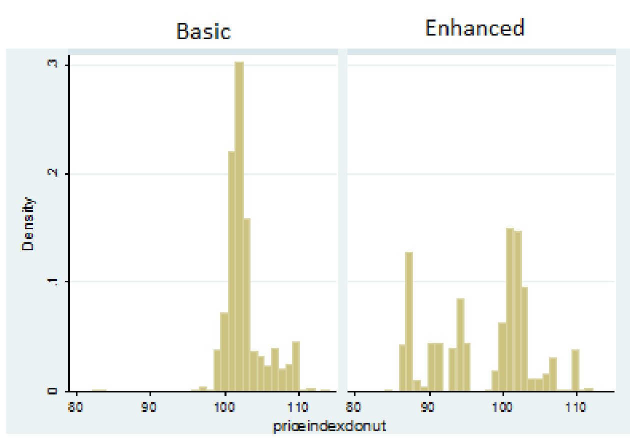

Figure 9: **2009: Donut Hole Price Index Histogram**
This histogram compares the donut hole out of pocket price index of basic and enhanced plans. Notice very few basic plans have an index below \$100, while many enhanced plans are below \$100. These enhanced plans provide some coverage in the donut hole.