Risk Adjustment and Low Income Subsidy Distortions in Medicare Part D

Daniel P. Miller*
Clemson University
Working Paper

June 27, 2015

Abstract

Many health insurance programs in the United States have been shifting towards private health insurance exchanges to harness the benefits of market competition. Designers and regulators of the programs face a challenge to preserve competitive incentives while also complying with legislative rules mandating that the exchanges serve a variety of socioeconomic populations with a wide range of health risks. This paper reveals a set of distortions in the largest health exchange, the Medicare Part D prescription drug program, related to the design of the Low Income Subsidy (LIS) program and the three "Rs" of Part D's risk sharing mechanism (risk adjustments, risk corridors, reinsurance). I document price distortions, biases in risk sharing payments, and evidence of insurers and drug suppliers using sophisticated drug price discrimination practices to exploit these distortions and biases. In conclusion, I discuss several policy considerations for designing health exchanges.

1 Introduction

Health insurance programs in the United States have been shifting towards a reliance on private insurance markets to harness the benefits of market competition. Large programs include Medicare Part D, Medicare Advantage, Medicaid Managed Care, and the new health insurance exchanges created under the 2010 "Patient Protection and Affordable Care Act." Legislative requirements dictate that these markets include subsidies to guarantee affordability —typically more generous for low income individuals—and prohibitions against experience rating to ensure equal access to insurance regardless of health status. To meet

*Author correspondence: dmille7@clemson.edu Financial support and data access provided through grant: "Low Income Subsidy Research Program" Patient Centered Outcome Research: U.S. Department of Health and Human Services Office of the Assistant Secretary for Planning and Evaluation.

these requirements, designers and regulators of the exchanges apply quite complicated and complex subsidies rules and risk sharing mechanisms. However, subsidies and risk sharing can act to blunt and distort market incentives, potentially undermining the objectives of the program. This poses a great challenge particularly when the rules and regulations become layered with complexities.

In this paper, I investigate the subsidies and risk sharing mechanisms in the Medicare Part D prescription drug program. Part D is the largest health insurance exchange covering over 41 million individuals. It has been operational since 2006. I focus specifically on the Low Income Subsidy (LIS) program that targets low income Medicare beneficiaries and the three "Rs" of the risk adjustment mechanism (risk adjustments, risk corridors, reinsurance). LIS beneficiaries compose about 20% of the Medicare population and a third of enrollees in Part D drug plans. As a group, they have some of the most severe health risks in the population and consume a disproportionate share of prescription drugs.

Despite well-thought intentions, the subsidy rules decouple plan choice from market fundamentals, which, combined with an imperfect risk sharing mechanism, acts to distort market outcomes. As an example, there is a particularly dubious symptom that defies economic intuition. Twelve percent of insurers' enhanced—so called "Cadillac"—plans are priced lower than that insurer's corresponding basic coverage plan. This headliner distortion is not necessarily the most significant, but it is endemic of the more critical, root causes. The problems may be manifesting in more severe forms that affect drug prices. Congress has taken notice. A recent investigation has shed light on skyrocketing prices of generic drugs.[1] Congressional leaders recently proposed a bill that would institute greater government price controls.[2] In the face of the concerns about rising drug prices, Part D has been experiencing a paradoxically trend. Premiums (gross of subsidies) are rapidly falling, while the government's share of expenditures on Part D are rapidly rising. The results of this paper suggest a mechanism linked to Part D causing the rise drug prices and divergent trends in program cost. The findings have implications about how anti-trust authorities should scrutinize the rapid wave of horizontal mergers amongst insurers and vertical integration of insurers with pharmacy benefit managers (PBMs) and pharmacy outlets.

To investigate these distortions, I specify a demand-side and supply-side model of the Part

[1] October 2, 2014 U.S. Sen. Bernie Sanders (I-Vt.) and U.S. Rep. Elijah E. Cummings (D-Md.) launch an investigation into soaring generic drug prices. Hearing on November 20, 2014 proposes the "Medicaid Generic Drug Price Fairness Act" to extend Medicaid drug rebates to generic drugs.

[2] On April 23, 2015 the "Medicare Drug Savings Act of 2015" (H.R. 2005 and S. 1083) was proposed to require drug manufacturers to provide drug rebates for drugs dispensed to low-income individuals under the Medicare prescription drug benefit.

D market including rich details about the LIS program, the three "Rs" of the risk sharing mechanism, and the interaction of subsidies and risk sharing. Using both publicly available data on Part D and restricted-access administrative claims data, I document distortions to enrollees' plan selection patterns, insurers' pricing and plan offerings, biases in the risk sharing mechanism, and evidence of risk sharing blunting insurers' incentives to control cost. In conclusion, I discuss a variety of policy options.

The remainder of the paper is organized as follows. In section 2, I introduce the details of the LIS program and risk adjustment mechanism to provide intuition on the distortions. In section 3, I specify and estimate a discrete choice demand model that shows how the subsidy rules affect consumer plan selection. In section 4, I present a supply-side model under a baseline assumption of perfect risk adjustments to show how subsidies distort insurer pricing and plan offerings. In section 5, I specify a model of the risk adjustment mechanism, present evidence as to how and why it is biased, and relate that bias back to the distortions in insurer pricing and plan offerings. In section 6, I model reinsurance and present evidence about how reinsurance payments are causing program costs and relate the rise to drug prices for LIS beneficiary claims. I use restricted-access administrative claims data on pharmacy transactions to show how these market distortions facilitate drug price discrimination that preys on the weakened incentives of insurers to control costs. Section 7 discusses policy considerations and conclude.

2 Low Income Subsidy (LIS) Program

Regular (non-low income) Medicare beneficiaries seeking coverage from the exchanges select plans based on price and coverage characteristics. Insurers set prices and negotiate drug price discounts with drug suppliers under a mild set of government constraints. When enrollees choose plans based on price and coverage, insurers have a strong incentive to compete on price and to control drug costs. Regular enrollees receive what is effectively a voucher subsidy that minimally distorts pricing.

Prior to Part D, low income Medicare beneficiaries dually qualified for Medicaid received drug coverage through Medicaid's drug program. When Part D was introduced, these "dual eligible" Medicare/Medicaid individuals lost their Medicaid drug coverage and were moved over to coverage being offered by private Part D insurers; the same set of plans offered to regular enrollees. For dual eligibles and other very low income individuals, legislation deemed it necessary to offer them affordable, in fact (nearly) free drug coverage, and to guarantee coverage for dual eligibles who lost Medicaid coverage. The law dons responsibility on the government to make plan assignment decisions on behalf of dual eligibles who do not choose a plan own their own accord. Mandates for free, guaranteed coverage in which the individual has no obligation to choose stifle incentives to compete. Designers of Part D faced the dual challenge of meeting these mandates while also creating a marketplace that encourages competition.

Part D includes special stipulations for low income and dual eligible 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 drug copays over and above those available to regular enrollees. For those at the lowest income levels, drug coverage can be essentially free. To provide competitive incentives, the rules stipulate that a plan is eligible for the full LIS subsidy if it prices below an endogenously determined threshold. Roughly speaking, the threshold is calculated as the average premium in the market. Low income beneficiaries selecting more expensive plans have to pay the premium difference on their own. Second, Medicare/Medicaid dual eligibles, who do not actively choose a plan are automatically and randomly assigned to a plan. About 77% of dual eligibles do not choose a plan and once auto-assigned only 10% subsequently turn down their auto-enrollment (Summer et al., 2010). Only plans priced below the threshold are eligible to receive auto-enrollees, creating an incentive for plans to compete for auto-enrollees and preventing them from being assigned to high cost plans. To ensure competition persists across years, plans that maintain full LIS eligibility have their auto-enrollees rolled over, and plans that lose eligibility lose their auto-enrollees who are then randomly reassigned to another plan.

Subsidies, automatic enrollment, the endogenous threshold, and year-to-year reassignment all act to distort insurers' decisions about pricing and plan offerings. I focus on two key results. First, automatic enrollment creates a demand discontinuity at the LIS threshold resulting in a bunching of prices at the threshold. I examine the question of whether the observed bunching promotes or hinders competition. Second, the endogenously determined threshold induces plans to cycle in-and-out of LIS eligibility (above-and-below the threshold) each year. This cycling causes dual eligibles to be assigned and re-assigned to plans year after year, which can have welfare consequences for dual eligibles because it disrupts coverage continuity. Related work in Decarolis (2015) and Ericson (2014) further explore insurers' strategic, in particular dynamic, pricing incentives.

2.1 The Three "Rs": Risk Adjustments, Risk Corridors, Reinsurance

The full explanation of the distortions caused by the LIS program hinges on another component of the Part D market design: the three "Rs" of the risk sharing mechanism (risk adjustments, risk corridors, and reinsurance). These risk sharing components were put into Part D to mitigate adverse selection (Glazer and McGuire, 2000) and to reduce risk for insurers entering a new and unknown market. Similar mechanisms can be found in the health exchanges created under the Affordable Care Act and Medicare Advantage. The connection to the LIS program not only nuances how one should interpret the subsidy distortions described above, but also sheds light on a set of issues that have puzzled regulators (CMS: Center for Medicare/Medicaid Services) since the program's inception. Transfer payments between plans and the government for the three "Rs" have never come in as expected and exhibit great variability. I propose a theory to explain the discrepancies and present evidence of market forces linked to the LIS program that are systematically biasing the risk sharing mechanism.

Risk adjustments, the first "R", are transfer payments between plans. Transfer amounts are based on measurable chronic condition risk factors and demographics of a plan's enrolled patient pool. For example, a plan with a 10% "sicker" than average pool of enrollees receives 10% extra payment. Conversely, a plan with 10% "healthier" than average enrollees has 10% of its payment deducted. In principal, insurers should be indifferent as to whether they enrollee a high or low risk pool of patients, which should help alleviate adverse selection problems. However, the regression models used to predict risk are not perfect predictors of cost (low r-square), which presents opportunities for insurers to "cream skim" favorable risks. That is, they try to attract enrollees with expected drug expenditures below that predicted by the risk adjustment model and deter those with expenditures above. As the literature suggests (McAdams and Schwarz, 2007; Hsu et al., 2009; Carey, 2014), insurers adjust their benefit design to cream skim using sophisticated formulary management techniques that target specific health conditions.

There is a key connection between risk adjustments, cream-skimming, and the LIS automatic enrollment provision that boils down to the matter of choice. Insurers' cream-skimming techniques are only effective for enrollees actively choosing plans based on price and coverage features, not for dual eligibles randomly assigned to plans. An analogous concept has emerged in the health insurance literature regarding the interaction of risk selection and switching costs (Handel, 2013; Polyakova, 2015). Enrollees facing high switching costs are not actively choosing plans just as auto-enrollees are not actively choosing plans. The discord implies that (non-cream skimmed) auto-enrollees are more costly on a risk adjusted basis than (cream skimmed) regular enrollees.

As a result of this cost difference, there is not only a demand discontinuity at the LIS threshold but also a cost discontinuity. The combination of a demand and cost discontinuity results in a bunching of prices below *and above* the threshold. Bunching above the threshold is clearly bad for the market because it effectively turns the LIS threshold into a price floor. Insurers would like to price lower to attract more regular enrollees, but cannot because they want to avoid LIS beneficiaries. In the data, bunching above is clearly evident through careful inspection of pricing around the threshold. There are distinct modes in the distribution of pricing just above, and just below the threshold, with a gap in between that can be attributed to insurers' imperfect information about the location of the threshold. The large number of plans priced just-above was first noticed by CMS in the second year of the program (2007). CMS then issued the *DeMinimis* stop-gap rule that "softened" the threshold so that plans priced within a couple dollars above the threshold could retain LIS eligibility. In light of the theory and evidence, I argue that the deminimis rule has had little effect.

Risk corridors, the second "R", is a profit/loss risk sharing scheme between insurers and the government. Under risk corridors, insurers are compensated by the government if their realized costs are higher than that predicted by the risk adjustment formula. Likewise, the government deducts payments from plans that have realized costs below that predicted by the risk adjustment formula. The original purpose of risk corridors was to insure insurance companies against aggregate cost shocks in a new and unknown market. The downside to risk sharing with the government is that it blunts insurers' incentives to control cost, in particular their ability to negotiate low drug prices. This could potentially undermine one of the key tenants of Part D to rely on insurers to negotiate drug prices with little government interference. Throughout the program's history, realized costs have deviated significantly from the risk adjustment formula's predicted costs and exhibit a high degree of variability across insurers and volatility for any given insurer across years. The deviations and volatility have prompted CMS to indefinitely continue risk corridors despite the original legislative intent to discontinue risk corridors in 2008.

I present theory and evidence from a differences-in-differences model applied to risk sharing data to show that the deviations and volatility are a systematic result of the cost differences between regular and LIS beneficiaries stemming from cream-skimming. As insurers' cycle their plans in and out of LIS eligibility across years, their risk corridor payments systematically cycle up and down. As originally intended, there is no harm in reducing or possibly eliminating risk corridors.

Reinsurance, the third "R," is dollar-for-dollar perhaps the most important component of the risk sharing mechanism. Under reinsurance, the government bears the majority of insurance risk for enrollees with very large "catastrophic" drug expenditures, exceeding $5100 annually (in 2006). Below catastrophic levels, insurers and enrollees are responsible for expenditures. The government pays the enrollee share for LIS beneficiaries. Reinsurance composes a large fraction of total government expenditures on the Part D program and has grown at a faster rate than any other budgetary component of the program: doubling from $7 billion in the initial years to $14 billion in 2012 and projected at $24 billion for 2014. LIS enrollees compose a disproportionate share of the population covered by reinsurance. Only 3% of regular enrollees reach reinsurance spending levels, 20% for LIS enrollees.

I show that the same systematic biases in risk corridor payments appear in reinsurance payments. However, the bias is larger—too large—upwards of 3 billion dollars. I present evidence that could explain the excess bias based on drug price discrimination practices. I use restricted-access administrative drug claims data to show that drug prices are higher when the claim is covered by the government through either reinsurance or direct subsidies to LIS beneficiaries. Part D rules expressly prohibit insurers and drug suppliers from price discriminating against individual claims. However they are able to price discriminate against LIS and reinsurance claims using more sophisticated second and third degree techniques targeting drug formulations, pharmacies networks, and seasonality. The LIS automatic enrollment provision further facilitates price discrimination by separating regular and low income beneficiaries into different plans.

The drug price discrimination results showcase the cost-controlling incentive problems of government risk sharing. The high level of risk sharing via risk corridors could also be contributing to high drug prices. The problem extends beyond just drug prices. The price discrimination is a further detriment to the program because the budget-neutrality triggers in the legislation reduce the (less-distortionary) direct subsidies for premiums as reinsurance payments rise.

The are several policy approaches to remedy the distortions. CMS has already taken action. In 2011, the risk adjustment formula was revised and includes two major changes. First, risk adjustments are calibrated to drug expenditure data from prior years of Part D experience. At the outset, there was no Part D expenditure data. The formula in force from 2006 to 2010 was calibrated to an out-of-market sample of Medicare beneficiaries in the Federal Employee Health Benefits program and Medicaid drug programs. Second, risk adjustments are separately calculated for the LIS and regular segments of the market. Statistically, this should eliminate discrepancies in risk adjustment payments between the two segments. I present evidence that it may have worked to achieve this aim, but there is still a high degree of volatility and imbalance.

I question whether these two revisions are the best of course of action. Instead they may be putting the program on a path to insolvency. The problem is that risk adjustment payments become endogenous to the market. Higher drug prices this year for LIS beneficiary claims lead to higher risk adjustment payments for LIS beneficiaries in the future. The prospects of higher risk adjustment payments in the future heighten drug suppliers' incentives to price discriminate. These incentives create a cycle of rising drug prices that increase government payments on reinsurance and LIS claims, which erodes the otherwise well-designed market mechanisms that promote competition. This cycle may have just reached an inflection point. 2014 was the first year that risk adjustment payments were based on drug prices after the 2011 revision. Congressional leaders have taken notice of rising drug prices and proposed price controls targeting claims for LIS beneficiaries and even generic drugs.[3] In conclusion, I offer several alternative policy thoughts.

3 Demand Model

I model demand for plans using the discrete choice framework of Berry (1994); Berry et al. (1995).

3.1 Demand: Regular (non-Low Income) Enrollees

Every year t, a consumer, indexed by i, can enroll in one prescription drug plan. Consumers choose from amongst the $j = 1, \ldots, J_{mt}$ differentiated plans offered in market m in year t. Markets are geographically separated into 34 regions drawn around state borders. They may also choose an outside option, $j = 0$ with utility normalized to zero. The outside option includes foregoing drug coverage, enrolling in a bundled MA+Part D plan, or getting coverage from another source, such as a current employer, another government program, or a Retiree Drug Subsidy (RDS) program plan.

Enrollees pay a premium p_{jmt} set by the plan. They derive utility from plan characteristics

[3] See reference to bills in introduction section.

and income left over after paying the premium. Define the conditional indirect utility of consumer 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, including coverage measures such as the deductible, drug copay/coinsurance rates, and the size of pharmacy networks. The term ξ_{jmt} represents an index of unobservable (to the econometrician) plan characteristics, including such non-fiduciary plan attributes as marketing activities, customer service qualities, and claims processing reliability. The terms ϵ_{ijmt} capture idiosyncratic differences in consumers' preferences over plans, which I interpret as match values between a patient's drug regimen and a plan's formulary composition/restrictions/pricing tiers over the set of 5000+ Part D drugs. The terms α_i, and β_i are random coefficients that represent consumer i's marginal utility over income and plan 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 they perceive to yield the highest conditional indirect utility in equation 1.

3.2 Demand: Low Income Subsidy Enrollees

The utility specification in equation 1 can be explicitly modified to account for the features of the low income subsidy program: premium subsidies, drug cost sharing reductions, and the automatic enrollment provision for dual eligible beneficiaries that do not actively select a plan.

An enrollee's eligibility for the low income subsidy is described by the parameter $\kappa_i \in [0, 1]$. An 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 Medicare beneficiaries that are enrolled in Medicaid—dual eligibles—are automatically granted eligibility at $\kappa_i = 1$. Second, non-Medicaid beneficiaries 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 $\kappa_i = 1$. The subsidy parameter decreases in increments of 0.25 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.

The low income premium subsidy is a function of the plan's coverage designation, its own premium, and premiums set by other plans in its market. Plans are designated as either basic plans or enhanced plans. Basic plans meet (or are actuarially equivalent) to the minimum coverage standards in the Part D legislation. Enhanced plans offer coverage exceeding the minimum standards, typically by lowering the deductible, reducing drug copay/coinsurance rates, or by covering non-Part D drugs. A plan's total premium, p_{jmt}, is calculated as the sum of a basic premium component, p_{jmt}^{basic}, and an enhanced component $p_{jmt}^{enhanced}$:

$$p_{jmt} = p_{jmt}^{basic} + p_{jmt}^{enhanced}. \tag{2}$$

Only the component of the premium attributable to basic coverage, p_{jmt}^{basic}, is subsidized. The subsidy amount is capped at a threshold, $\bar{s}_{mt}^{LIS}$, determined by taking a weighted average of the basic components of pricing for all plans in the market. The component of the premium attributable to enhanced coverage, $p_{jmt}^{enhanced}$, is not subsidized. I further elaborate on the threshold rules and distinction between basic and enhanced components in the supply-side discussion. The *full* low income premium subsidy amount in market m for plan j, $\tilde{s}_{jmt}^{LIS}$, is the lesser of the plan's basic premium or the market threshold:

$$\tilde{s}_{jmt}^{LIS} = min\{p_{jmt}^{basic}, \bar{s}_{mt}^{LIS}\}. \tag{3}$$

The subsidy received by an enrollee of type κ_i is $\kappa_i \tilde{s}_{jmt}^{LIS}$. Figure 1 illustrates the after subsidy, out-of-pocket premium $(p_{jmt} - \kappa_i \tilde{s}_{jmt}^{LIS})$ for enrollees of various levels of κ_i 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 shifts the payment schedule. Note that it is possible (and observed in the data) for an enhanced plan to have a lower total premium p_{jmt} than a basic plan, yet have a higher after subsidy premium.[4] This subsidy anomaly is part of explanation for the teaser fact from the introduction about enhanced plans being priced lower than basic plans.

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

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

[4] For example, a \$30 enhanced plan with basic/enhanced components (25/5) would have a higher after subsidy price (5) than a \$35 basic plan in a market with a threshold of \$35.

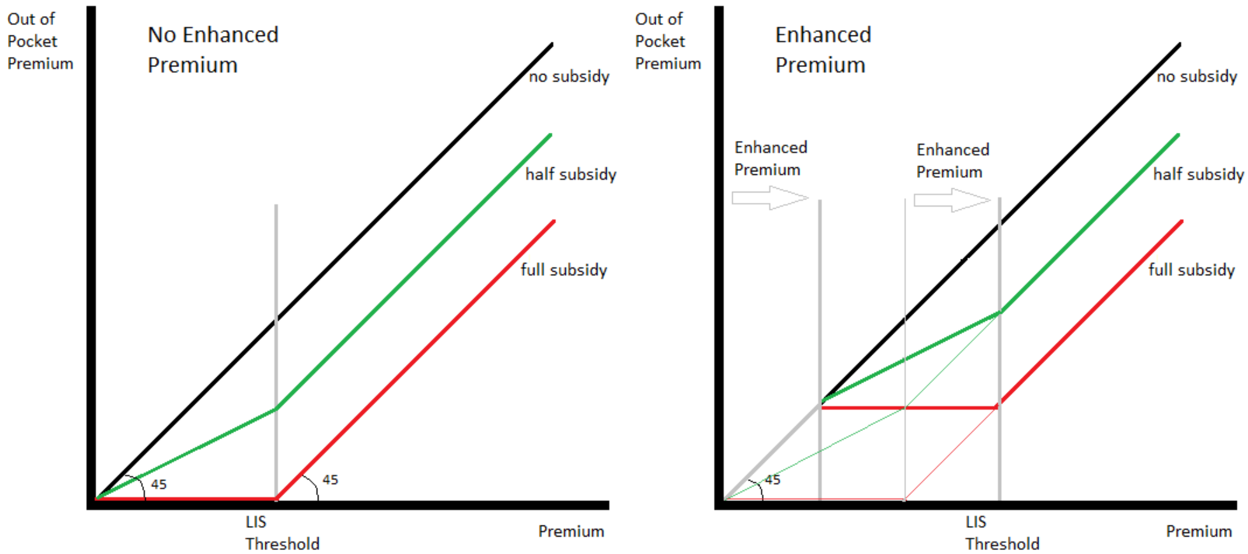

Figure 1: Low Income Subsidy Out-of-Pocket Premium

Under this specification, enrollees receive disutility from choosing a high premium plan and gain utility from choosing a plan with a high subsidy. The specification explicitly models the non-linear relationship between premiums and subsidies and nests the generic utility specification in equation 1 if $\kappa_i = 0$.

At the market level, I observe separate enrollment figures for regular enrollees ($\kappa_i = 0$) and for the subset of the population that has LIS eligibility $\kappa_i > 0$. I estimate separate demand models for each segment of the population with no restrictions placed on the preference parameters $(\alpha_i, \beta_i, \xi_{jmt})$ across segments. For example, its plausible that the random coefficient distributions of α_i and β_i differ between regular and low income enrollees due to differences in income levels and health risks.

For regular $\kappa_i = 0$ enrollees, I impose the restriction that the subsidy amount does not affect utility, and use the baseline utility model in equation 1 to estimate demand.[5] I use the subsidy amount as an exclusion restriction to instrument for the premium.

For the LIS eligible segment of the market, I substitute the subsidy amount into utility according to equation 5. CMS restricts information about the distribution of enrollee types κ_i. I estimate the following restricted version of the model,

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}^{LIS}_{jmt}; \kappa_i) = -\alpha_i p_{jmt} + \alpha^s_i \tilde{s}^{LIS}_{jmt} + \mathbf{X}'_{jmt}\beta_i + \xi_{jmt} + \epsilon_{ijmt} \tag{5}$$

[5]There may be some marketing value for a plan to advertise that it is zero premium eligible. But such marketing is targeted towards low income beneficiaries, not regular enrollees.

where $\alpha_i^s = \kappa_i \alpha_i$ is a distinct random coefficient from α_i.

I make assumptions about the joint distribution of α_i and α_i^s to estimate the model for the subset of LIS eligible households. I begin by assuming the distribution of α_i is normal: $\alpha_i \sim N(\bar{\alpha}, \sigma^2)$. There are three points to make about the distribution of α_i^s. The first, most restrictive assumption I could make assumes $\kappa_i = \kappa$ for all i. In this case, 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, more flexible assumption I could make does not restrict κ_i to be fixed but instead 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. I would expect positive correlation between α_i and α_i^s if there is diminishing marginal utility of income over the relevant income range below 150% of the federal poverty line.

I 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, I do not expect this to be too restrictive because a very large mass of households have $\kappa_i = 1$. The range between 135% and 150% of the federal poverty line is quite narrow for those qualifying through means testing. Nationwide statistics published by CMS indicate about 75% of LIS recipients qualify as $\kappa_i = 1$ dual eligibles. As an approximation, restricting $\alpha_i = \alpha_i^s$ would be appropriate if all LIS recipients have full eligibility. The multivariate normal distribution allows for some departure from this strict restrictions.

The low income subsidy program also reduces the deductible amount and copays/coinsurance rates in the initial coverage zone and donut hole. Like premium subsidies, copay subsidies depend on κ_i. The deductible is \$0 for all $\kappa_i = 1$ enrollees. For 2009, the maximum deductible is capped at \$60 for $\kappa_i \in (0,1)$ and \$295 for $\kappa_i = 1$. These caps increase each year. The initial coverage zone is the range of drug expenditures after the deductible has been met and before the so-called "donut hole" gap in coverage (\$295 to \$2900 in 2009). In the initial coverage zone, the coinsurance rate for basic plans should be actuarially equivalent to 25% for regular $\kappa_i = 0$ enrollees. All LIS recipients with $\kappa_i \in (0,1)$ have a coinsurance rate capped at 15%, and for fully eligible $\kappa_i = 1$ individuals cost sharing is set to a nominal copay of \$2.40/\$6.00 (generic/branded). In the donut hole (\$2900 to \$6150[6] in 2009), regular

[6]The upper limit on the donut hole is based on a maximum of out-of-pocket expenditure of \$4350 which translates to approximately \$6150 in drug expenditures.

enrollees receive no cost sharing benefit. The donut hole is eliminated for all enrollees with $\kappa_i > 0$. These levels represent maximum subsidy amounts. Deductibles and cost sharing may be more generous than these levels for individuals purchasing enhanced plans. For example, many enhanced plans have a \$0 deductible and coverage in the donut hole.

Finally, I must account for automatic enrollment of dual eligible beneficiaries. Medicare randomly assigns all dual eligibles who do not actively enroll in the plan. 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 to accept random assignment. At anytime, even outside the open enrollment period, an auto-enrolled beneficiary is allowed to choose another option.

From the perspective of a dual eligible, neglecting to enroll and accepting random assignment can be an attractive option because the plan will have zero premium, zero deductible, and near zero copays/coinsurance. At first glance, dual eligibles should be indifferent amongst all fully eligible LIS plans. But consider utility with no premium or cost sharing,

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}^{LIS}_{jmt}; \kappa_i = 1) = \xi_{jmt} + \epsilon_{ijmt}. \tag{6}$$

There are two terms in the utility function: unobserved plan characteristics, ξ_{jmt}, and the idiosyncratic preference shock, ϵ_{ijmt}. The terms reflect non-fiduciary plan qualities and idiosyncratic characteristics about the composition and restrictions of the plan's drug formulary. All enrollees, regular and low income, must adhere to a formulary's drug exclusions and usage restrictions. Given there is idiosyncratic heterogeneity in patients' drug regimens, the ϵ_{ijmt} terms represent a match value of patient to formulary.

For auto-enrollees there exists some plan that is a best match. The random nature of auto-assignment does not guarantee assignment to the best plan. There are several reasons in the consumer choice (Klemperer, 1995) and Part D literature (Ketcham et al., 2012; Kling et al., 2012) suggesting why beneficiaries prefer accepting the default auto-assigned plan. Perhaps most important are inattention biases (Madrian and Shea, 2001) and the time and effort costs of researching formulary details to forecast which plan is the best. If these cost exceed the perceived difference in utility between the randomly assigned plan and best alternative, the enrollee will accept random assignment. Severe mismatch would induce the selection of an alternative. Much of the literature on plan choice (Abaluck and Gruber, 2011; Heiss et al., 2013) documents behavioral irregularities for regular enrollees, and the same principles likely apply for dual eligibles regarding their decision to opt out of auto-enrollment.

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

$$U_i(X_{jmt}, p_{jmt}, \tilde{s}^{LIS}_{jmt}; \kappa_i) = -\alpha_i p_{jmt} + \alpha^s_i \tilde{s}^{LIS}_{jmt} + \mathbf{X}'_{jmt}\beta_i + \beta^{lis}\mathbf{1}(LISPLAN_{jmt}) + \xi_{jmt} + \epsilon_{ijmt} \quad (7)$$

The coefficient β^{lis} determines the proportion of households that are randomly assigned. In utility terms, it is the value low income beneficiaries place on being auto-assigned. Holding all else fixed in the utility specification, the market shares amongst LIS plans induced by the logit model distributes automatic enrollees uniformly across LIS eligible plans. CMS applies a uniformly random assigned process, with exceptions in a few markets.[7] The idiosyncratic error terms, ϵ_{ijmt}, take on an additional interpretation as the random number generator determining which plan an auto-enrollee gets assigned to.

This is a static model of auto-enrollment that abstracts away from some additional rules linked to dynamic considerations. I consider dynamic issues on the supply-side. From year to year, auto-enrollees are kept in the same plan if it maintains its LIS qualification. If the plan loses its qualifications, the auto-enrollees are redistributed uniformly amongst other LIS eligible plans. The *DeMinimis* rules allow incumbent LIS eligible plans priced slightly above the threshold to retain auto-enrollees, however they cannot receive new auto-enrollees. Finally, a dual eligible loses all future rights to be auto-assigned upon the first occasion that he actively selects a plan. Dual eligibles that actively select are called "choosers."

3.3 Demand Estimation Results

Table 1 presents demand estimates for the LIS and non-low income, regular segments of the market using data from 2009. The product characteristics include the monthly premium p_{jmt}, LIS subsidy $\kappa\alpha$, and LIS plan dummy variable discussed above. Coverage characteristics include the annual deductible divided by 12, and measures of monthly drug copay/coinsurance rates in the initial coverage zone and donut hole. The copay/coinsurance rate variables are constructed as a price index of the out-of-pocket price a regular (non-LIS) enrollee would pay at a network pharmacy to fill a basket of the top 100 most popular drugs under a plan's specific copay/coinsurance cost sharing rules. I set the price of drugs excluded from the

[7] A few states, such as Maine, have adopted experimental assignment programs that attempt to match auto-enrollees based on drug regimens and formularies(Zhang et al., 2014).

formulary to the full pharmacy retail price. There is variation across plans in out-of-pocket prices stemming from differences in negotiated drug prices, formulary exclusions, cost sharing tiers (preferred/non-preferred), and coverage enhancements for enhanced plans. The final characteristic is a count of the number of in-network pharmacies contracting with the plan. Instruments for the endogenous premium and LIS subsidy variable include BLP instruments measuring isolation in product space and exclusion restrictions implied the regulatory framework. The LIS plan indicator variable and LIS subsidy variables are used as instrument for the regular population which is not subject to auto-enrollment and receives no LIS subsidy. The donut hole price index variable instruments for the LIS population which does have a donut hole gap in coverage.

The first two columns report results for a non-random coefficient logit specification. For both the LIS and regular population, demand increases as premiums fall and coverage characteristics improve. For the LIS population the coefficient on the LIS subsidy is slightly lower in magnitude than that for the premium because some LIS enrollees have $\kappa_i < 1$. They would be equal if $\kappa_i = 1$ for all LIS beneficiaries. The coefficients on coverage characteristics are similar for the regular and LIS segments. The similarities reflects a balance of the LIS population having truncated sensitivities to coverage because of subsidies, yet having a stronger preferences for coverage because they have very high levels of drug expenditures.

The second two columns report results for the random coefficient specification. The correlation coefficients amongst all random coefficients on the premium, LIS subsidy, and out-of-pocket drug price indices are fixed at a value of 0.985. The high correlation is sensible because all characteristics represent dollar-valued expenditures. For example, a person with a high marginal utility of income (α_i) also has a high marginal utility from paying the deductible. [8] Compared to the logit specifications, the mean value of the premium and LIS subsidy coefficients are much higher. The large standard deviation indicates significant heterogeneity in the population. Like the logit specification, the LIS subsidy coefficient is slightly lower than that for the premium because some beneficiaries have $\kappa_i > 0$. The coefficient on the premium for the low income LIS population is double that of regular enrollees, which is indicative of diminishing marginal utility of income.

The coefficient on the LIS indicator variable determines how much demand shifts for LIS eligible plans due to auto-enrollment. There are two ways of interpreting the coefficient in the context of the logit model. First, it can be interpreted from a utility perspective.

[8] I experimented with unrestricted parameterizations of correlation coefficients that converged towards correlation coefficients of 1, but there are numerical stability issues with the BLP algorithm near the boundary. I found the best GMM objective function fit at 0.985.

Table 1: Demand Estimation Results

	LIS population Logit	regular population Logit	LIS population RC	regular 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 random coefficients=0.985

Auto-enrollment improves the utility for LIS beneficiaries because it removes the burden of having to choose a plan. The dollar value of the utility of auto-enrollment is $15, calculated by dividing through by the mean coefficient on α_i. In other words, a LIS beneficiary would override auto-enrollment if he could save $15 dollar per month ($180 per year). This is a substantial amount considering these individuals have incomes at very low poverty levels.

An alternative interpretation is to consider how auto-enrollment affects market shares. Consider three very different states. At one extreme, Nevada had 1 LIS eligible plan out of 49 total plans for 2009. Assuming mean utility, net of the LIS coefficient, is equal for all 49 plans, the single LIS plan would have a predicted market share of 35%. At the other extreme, South Carolina had 15 LIS eligible plans out of 53. Each LIS plan would have a market share of 6%, for a combined market share amongst the 15 of 90%. In the middle is New York with 51 plans in which each of its 9 LIS plans would have a 9% share, combined share of 81%. The points to notice are that the auto-enrollee share per plan decreases in the number of LIS plans because enrollees are uniformly distributed across plans. Also, the combined share of enrollees in LIS eligible plans increases as the ratio of LIS plans to total number of plans increases.

4 Supply with Perfect Risk Adjustments

I model the supply side by closely following the regulations in the Medicare Modernization Act of 2003 and subsequent reforms. This section describes the pricing and subsidy rules abstracting away from issues of risk adjustments and cream-skimming by assuming perfect risk adjustments. The model illustrates the basic intuition of how insurers respond to the LIS rules. The next section introduces the details of the risk adjustment mechanism.

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. The payment is risk adjusted based on disease and demographics. For now, I assume the payment is perfectly risk adjusted to reflect differences in cost across enrollee pools with differing risk factors. Part of that payment is made by enrollees in the form of the premium p_{jmt}, and the remainder is subsidized by the government.

I model a plan's (risk adjusted) marginal cost mc_{jmt} of enrolling an individual as a constant. The marginal cost can be separated into a basic and enhanced component.

$$mc_{jmt} = mc^{basic}_{jmt} + mc^{enhanced}_{jmt} \tag{8}$$

with an enhancement ratio defined as

$$\gamma_{jmt} = \frac{mc^{enhanced}_{jmt}}{mc^{basic}_{jmt}}. \tag{9}$$

By regulation $\gamma_{jmt} = 0$ for basic plans and is positive for enhanced plans. Enhanced plans with the most generous cost sharing provisions, such as eliminated deductibles and donut holes, have large values of γ_{jmt}.

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

$$\Pi_{ft} = \sum_{mt} M_{mt} \sum_{J_{fmt}} (b_{jmt} - mc_{jmt}) s_{jmt}(\mathbf{b}) \tag{10}$$

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. Market shares s_{jmt} are the sum of the demand for both regular and LIS beneficiaries, written to explicitly depend on the bid vector $\mathbf{b}$ for all plans across all markets.

4.1 General and Low Income Premium Subsidy Rules

The regulations set the rules for determining the size of the general premium subsidy (also called the direct subsidy) which applies for all enrollees, and the low income subsidy which only applies for LIS enrollees. The general premium subsidy is calculated as a fixed proportion, λ_t, of the enrollment weighted *average basic* bid component of all plans in the country ($\lambda_t \approx 65\%$). Later, I discuss how λ_t is determined, as it is linked to reinsurance. Like the premium, a bid b_{jmt} separates into a basic component b^{basic}_{jmt} and an enhanced component $b^{enhanced}_{jmt}$:

$$b_{jmt} = b^{basic}_{jmt} + b^{enhanced}_{jmt}. \tag{11}$$

The general subsidy, $\bar{s}^g_{jmt}$, for plan j in market m in year t is

$$\bar{s}^g_{jmt} = min\{b^{basic}_{jmt}, \lambda_t \bar{b}_t\}. \tag{12}$$

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

$$\bar{b}_t = \sum_{jt} \tilde{w}_{jt-1} b_{jt}^{basic}. \tag{13}$$

The weights $\tilde{w}_{jt-1}$ are based on the previous year's total enrollment E_{jt-1} including both regular and LIS enrollees,

$$\tilde{w}_{jt-1} = \frac{E_{jt-1}}{\sum_{jt} E_{jt-1}} \tag{14}$$

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.[9]

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

$$\begin{aligned} p_{jmt} &= p_{jmt}^{enhanced} + p_{jmt}^{basic} && (15) \\ p_{jmt}^{enhanced} &= b_{jmt}^{enhanced} && (16) \\ p_{jmt}^{basic} &= b_{jmt}^{basic} - \bar{s}_{jmt}^{g} && (17) \end{aligned}$$

Note that the general subsidy is capped by the basic component of the bid (equation 12) to prevent the basic component of the premium from being negative. Strictly speaking, it has never been a binding constraint. However, bidding data shows that it nearly binds for a subset of enhanced plans. I discuss these plans after introducing risk adjustments. Ignoring this constraint, the subsidy rules give all enrollees the same general subsidy amount regardless of plan choice. Enrollees realize cost savings from choosing cheaper than average plans and pay extra to pick one that is more expensive. The enhanced component of the bid is not subsidized.

I consider two interpretations of the rules regarding how firms choose b_{jmt}^{basic} and $b_{jmt}^{enhanced}$. 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. The most strict interpretation assumes plans choose the total bid b_{jmt} but have no discretion about allocating between the basic and enhanced components. With the proportion based on cost, the ratio between the basic and enhanced component of the bid is the same as that

[9]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.

between the cost components:

$$\gamma_{jmt} := \frac{mc^{enhanced}_{jmt}}{mc^{basic}_{jmt}} = \frac{b^{enhanced}_{jmt}}{b^{basic}_{jmt}} \tag{18}$$

Later, I consider a less stringent interpretation that assumes insurers have some discretion over how to allocate amongst the basic and enhanced components of the bid. This simpler assumption eases estimation of the model as it reduces the number of prices chosen by insurers from two to one.

The demand section introduced most of the rules for the low income premium subsidy. A key component, not yet discussed, is the LIS threshold, $\bar{s}_{mt}$. It is computed similarly to the general premium subsidy with some important differences. The threshold is the enrollment weighted *average basic* component of the premium for all plans in a market.

$$\bar{s}^{LIS}_{mt} = \sum_{jmt} \tilde{w}^{lis}_{jmt-1} p^{basic}_{jmt} \tag{19}$$

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

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

The weight is zero for plans that are new entrants to the market. Like the general premium subsidy, the weights transitioned from a simple average to weighted average up through plan year 2008. The threshold is bounded below by the minimum premium of a plan that only offers basic coverage. In the program's 9 year history, this has only been a binding constraint once (Nevada, 2009).

Although quite similar in logic, there are key differences between the overall subsidy and LIS threshold calculations. The LIS threshold is calculated at the market level, not national level; it only considers LIS enrollment, not total enrollment; it is based on the basic component of the premium, not basic component of the bid. These differences have important implications for firms' pricing strategies.

4.2 Pricing with Subsidy Distortions

The subsidy rules and LIS threshold distort firms' pricing decisions in quite complicated ways. I focus on the most salient distortion which can be illustrated with the aid of a simple supply and demand diagram. See the elasticity calculations in the appendix and

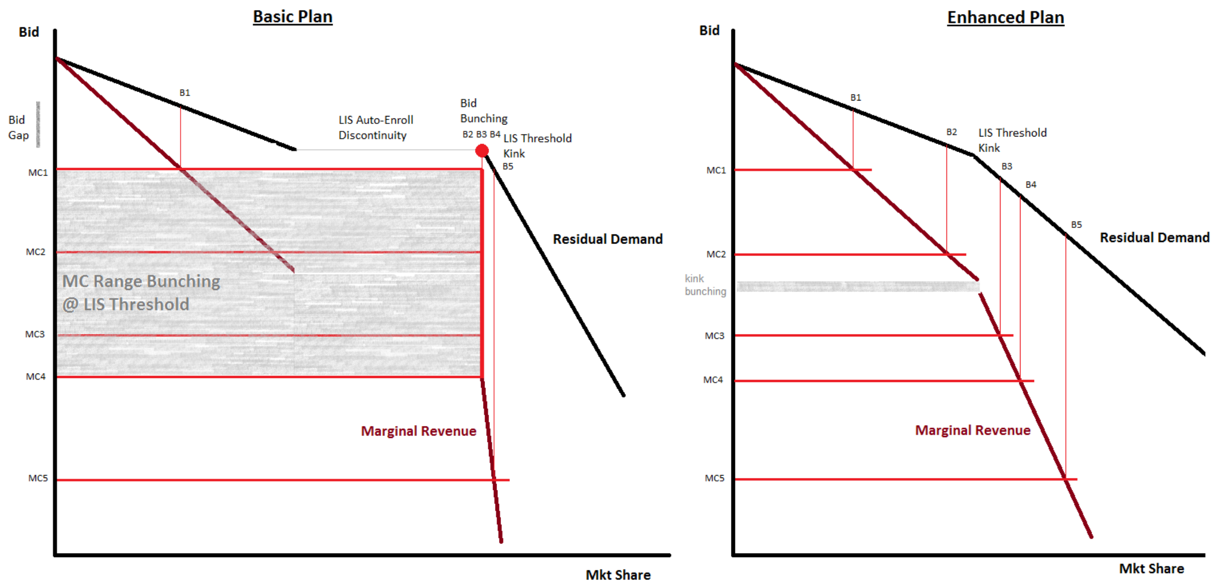

Figure 2: Demand Discontinuity

Decarolis (2015) for further discussion on pricing strategies. The LIS threshold creates a "discontinuity" and "kink" in residual demand as depicted in figure 2. The diagram depicts residual demand, marginal revenue, and cost curves. The first panel represents basic plans. There is a discontinuity and kink in demand at the LIS threshold. The second panel illustrates enhanced plans. There is a kink, but no discontinuity.

The first point to notice is the demand discontinuity for basic plans. For plans priced above the threshold, demand increases as the bid falls and then there is a large boost in demand at the LIS threshold because basic plans gain eligibility to receive LIS auto-enrollees. Enhanced plans do not have this discontinuity because they are not eligible to receive auto-enrollees. The size of the discontinuity depends on the "value" of auto-enrollment estimated in the demand model, β^{LIS}, and the number of rival plans priced below the threshold. From the prior example, the discontinuity is large in Nevada with 1 LIS eligible plan and small in South Carolina with 15 LIS plans. The discontinuity induces a bunching of prices at the threshold. As depicted in the figure, insurers with marginal cost in the grayed area (MC1-MC4), all want to price just below the threshold. Only firms with very low cost (MC5) would want to price lower than the threshold. There is a gap in pricing above the threshold. The intuition is that a plan would not want to price a few cents above the threshold because the large loss of auto-enrollees does not justify the incremental improvement in per-enrollee profit margins. A firm with cost at the top of the grayed area (MC1) would be indifferent between pricing at the threshold and the price point bounding the pricing gap. The size of the bid gap depends on the demand elasticity above the threshold and the size of the discontinuity.

The second point to notice is the kink at the LIS threshold. The kink is an artifact of how LIS recipients are subsidized up to the threshold, but not above. Residual demand is relatively elastic above the threshold and more inelastic below the threshold. How inelastic depends on the distribution of κ_i in the population. Beneficiaries with $\kappa_i = 1$ have perfectly inelastic demand for basic plans priced below the threshold, and beneficiaries with lower κ_i have progressively more elastic demand. The kink applies not only for basic plans but also enhanced plans because LIS subsidies also apply for the basic component of the premium for enhanced plans.[10] Like the demand discontinuity, the kink creates a discontinuity in marginal revenue curves. For basic plans, the kink amplifies the range of marginal costs that would price at the threshold. In theory, the kink induces a bunching of prices for enhanced plans, indicated by the grayed area. However, enhanced plan bunching is empirically quite small and negligible relative to that for basic plans because few LIS beneficiaries enroll in enhanced plans.

The pricing distortions caused by the LIS threshold are clearly evident in a kernel density plot of premiums. Figure 3 reports density plots of the basic component of the premium relative to the LIS threshold ($p^{basic}_{jmt} - \bar{s}_{mt}$) for basic plans, enhanced plans, and all plans pooled together in 2009. Plans to the left of the vertical line are priced below the threshold; plan to the right, above. The red density plots for basic plans shows a large mass of plans that set their premium right at, or just below the LIS threshold. There is a gap in the density just above the threshold, corresponding to the bid gap. A second mode appears well above the threshold, corresponding to the point labeled $b1$ in figure 2. By comparison, enhanced plans, marked in black, do not exhibit any bunching around the LIS threshold because they are not eligible to receive auto-enrollees. The mode for enhanced plans is above the LIS threshold in the region corresponding to the bid gap for basic plans. The comparison of basic and enhanced plans allays concerns that the spike for basic plans is purely an artifact of the underlying distribution of plans costs coincidentally coinciding with the LIS threshold. The grayed area is a kernel density plot for all plans pooled together. The second smaller mode for enhanced plans occurring well below the LIS threshold may be indicative of the binding lower bound on basic premium subsidies.

[10]Note that residual demand for enhanced plans is not perfectly inelastic below the threshold for $\kappa_i = 1$ beneficiaries due the enhancement ratio assumption on allocating the basic and enhanced component of the bid.

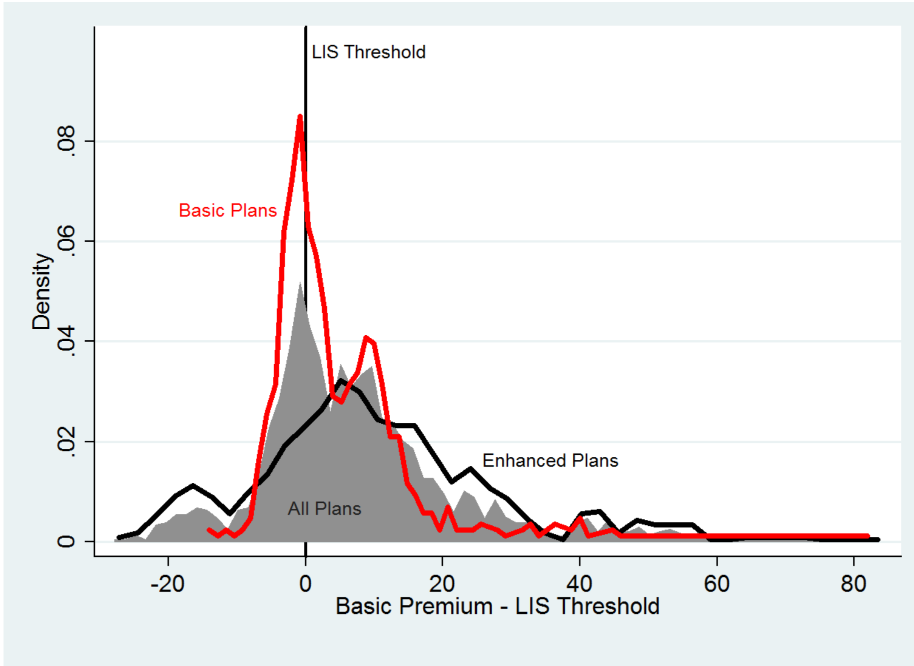

Figure 3: Bid Histogram: Bunching @ LIS Threshold

4.3 Does the LIS Threshold Rule Intensify or Soften Competition?

A natural question to ask is whether the LIS subsidy rules intensify or soften competition. On one hand, theory predicts softer competition for two reasons. First, insurers face very inelastic demand due to the generous subsidies for LIS recipients. Second, auto-enrollees are randomly assigned to plans with little regard for price or product characteristics. On the other hand, competition intensifies because the threshold rule requires plans to submit low bids to receive auto-enrollees.

A simple inspection of the pricing density plots in figure 3 suggests intensified competition. The basic component of the premium for basic plans appears lower than that for enhanced plans. The modal price of basic plans is at the threshold, about \$5 lower than the mode for enhanced plans. This is a large difference given the average monthly premium is about \$40. Because LIS subsidies and auto-enrollment affect basic plans and have little influence on enhanced plans, the comparison suggests the net effect of the LIS rules is to reduce prices.

Such a simple analysis makes implicit assumptions about competition and the distribution of cost across insurers. In particular, using market prices as a yardstick for competition assumes price differences are not driven by cost differences. A more detailed examination of prices for plans offered by the same insurer in the same market suggests reduced competition. Recall the teaser fact from the introduction. Of the 12% of insurers' enhanced plans priced below its corresponding basic plan, 35% of those basic plans are LIS plans. These insurers are offering a (slightly) enhanced plan priced low to attract regular enrollees and price higher at the LIS threshold to earn rents off of LIS auto-enrollees. Those auto-enrollees have no incentive to switch to the lower total premium $p_j^{basic} + p_j^{enhanced}$ enhanced plan because they would pay a positive dollar amount $p_j^{enhanced}$ for the slight enhancement (reduced deductible) which is already covered by the LIS subsidy.

To further gauge the question of whether the LIS rules soften or intensify competition, I conduct a structural estimation exercise to estimate cost and profit markups. The standard approach to solve for marginal cost in a Bertrand-Nash competition framework is to invert a system of first order conditions that maximizes profits in equation 10. The primary complication with this method is that the first order conditions do not hold with equality because of bunching at the LIS threshold. The mapping from bids to costs is not a one-to-one function because there is a range of costs mapping to a single price at the LIS threshold. I propose a straightforward method adapted from the assumptions of Bertrand-Nash competition to address bid bunching. Specifically, I use the first order conditions to estimate cost for enhanced plans, which hold with equality because there is no discontinuity in demand. Note, I disregard the discontinuity in marginal revenue due to the "kink" because it appears to have a negligible effect. Likewise, first order approaches apply for basic plans priced above and bounded away from the LIS threshold. To infer cost for a basic plan priced at (or very near) the threshold, I apply a restriction that it has the same basic component of the cost as its "sister" enhanced plan offered by same firm in the same market. The appendix describes details and limitations of the method.

I apply the method to one large national insurer in 2009 that has a mix of basic plans priced above and below the threshold. This particular insurer offers an interesting example that highlights how markups differ for basic plans priced above and below the threshold. The following tables report enrollment weighted averages for marginal cost mc_{jmt} and markups $(b_{jmt} - mc_{jmt})/b_{jmt}$. These are total costs and total bids, not separated into basic and enhanced components.[11] The first column applies the "sister" plan restriction. The second column, labeled "no-restriction," reports estimates under the assumption that the FOCs

[11] The results are based on the assumption that profits are proportionally allocated to the basic and enhanced component as implied by the strict enhancement ratio assumption in equation (18).

Table 2: Cost Estimates for a Large National Insurer

	"Sister" Restriction		No Restriction	
	MC ($)	Markup (%)	MC ($)	Markup (%)
Basic LIS	75.13	9.97	53.85	35.43
Basic Non-LIS	81.79	6.28	82.10	5.93
Enhanced	79.98	14.25	80.03	14.21

Marginal cost is reported as the total of the basic and enhanced components: $mc = mc^{basic} + mc^{enhanced}$. The Lerner markup is calculated as $\frac{b-mc}{b}$.

hold with equality for basic plans priced at the threshold. This naive approach is equivalent to assuming marginal cost is exactly at the lower bound of the cost range pricing at the threshold: MC4 as labeled in figure 2.

The results offer mixed conclusions about competition. Markups (10%) are higher for LIS plans priced below the threshold than for basic non-LIS plans priced above the threshold (6%). This comparison indicates auto-enrollment softens competition. Enhanced plans have very high markups (14%). Given enhanced plans are ineligible for auto-enrollees, a comparison to LIS plans indicates auto-enrollment intensifies competition. The naive, "no-restriction," estimation results would imply severely distorted competition: 35% markups for LIS plans.

The are various complications to estimate a structural model for the whole market. First is multiple equilibrium. The game played amongst insurers regarding decisions about whether to price above or at the threshold can be modeled as a discrete entry game. There is not necessarily a unique set of insurers who would price at the threshold because profits from auto-enrollees decline as more insurers enter at the threshold. As the entry literature shows, there is not necessarily a unique number of firms (Tamer, 2003). One can even envision more complicated representations of the demand curves in figure 2 with "stair-step" discontinuities in which an LIS firm would want to set a price bounded below the threshold to lower the threshold and "knock-out" a rival. There are further complications with the structural estimation approach regarding endogenous product positioning, imperfect information, dynamic pricing incentives, and risk selection. One of these anomalies appears in the result in table 2. Enhanced plans, with more generous coverage, have lower costs than non-LIS basic plans with less generous coverage. I elaborate on these issues in the next sections. For the purposes of this paper, I do not attempt a full blown structural estimation approach addressing all of these complications.

4.4 Endogenous Subsidy Dynamics: Plans Cycling In and Out of LIS Eligibility

In this section, I consider the effect that the endogenous subsidy has on insurers' pricing incentives, in particular dynamic pricing. The general subsidy and LIS subsidy are pegged to prior year (lagged) enrollment figures. Intuitively, small insurers with low lagged enrollment take both the general subsidy $\bar{s}^g$, and LIS subsidy amounts $\bar{s}^{LIS}$ as exogenous lump sum amounts that, on the margin, do not alter pricing decisions. Large insurers with high lagged enrollment influence the subsidy amount which they take into strategic consideration when setting prices.

Miller and Yeo (2013) consider the general subsidy and show that it creates more inelastic residual demand causing large firms to price higher. For example, an insurer with a 10% market share that raises the bid \$1 on its plans only raises its premiums by \$0.935 because the subsidy amount increases 6.5 cents (see equation 13). In the aggregate, the general subsidy distortion is a rather small amount, closely resembling a lump-sum voucher subsidy, because insurers have relatively low national market shares. The Hirschman Herfinhdahl Index (HHI) is a good proxy for gauging the distortionary effect. The 2014 national HHI is only 880, indicating the Part D market is "not-concentrated" according to Department of Justice guidelines. An alternative subsidy rule that would create a much larger pricing wedge would be to subsidize 65% of each plan's individual bid as opposed to 65% of the average bid. Decarolis et al. (2015) estimate a structural model of Part D and find that the general subsidy has a minimal welfare effect as compared to lump-sum voucher subsidy.

The same intuition regarding the pricing incentives of small and large insurers applies to the LIS threshold. However the distortionary effects are much greater because LIS thresholds are based on the local market enrollment of LIS recipients. In contrast, the general subsidy is determined by national enrollment of all enrollees. At the local level, the average HHI across the 34 markets for all enrollees is 1088. Only 4 of the 34 markets reach "moderately concentrated" levels between 1500 and 2500. However, the HHI measures of market concentration are much higher when they are based on LIS enrollment. LIS concentration reaches moderate levels in 11 of the 34 regions and crosses into the highly concentrated levels in 2 markets.[12] At high HHI levels insurers have more influence over the LIS threshold and subsidy, making it prone to manipulation.

The incentives of insurers to manipulate pricing have an inherent dynamic component

[12]Statistics on HHI are reported in Summer et al. (2014)

because the subsidy calculations depend on lagged enrollment. The dynamics are particularly important for the LIS threshold. The insurers with a large contemporaneous influence on the LIS threshold must have had high LIS enrollment in the previous year. Given the strong preference for auto-enrollment in the demand functions of LIS recipients, those insurers predominantly gain LIS enrollment by pricing at the LIS threshold. Once a plan is priced at the threshold, other market factors to gain LIS enrollment have little effect—such as further lowering the price or improving coverage desirability.

Decarolis (2015) proposes a model that shows an "unraveling" of the market that resembles predatory pricing. The lowest cost insurer (such as MC4 labeled in figure 2) lowers it bids to a point at which relatively high cost insurers (such as MC1 labeled in figure 2) drop out of LIS eligibility. The low cost insurer remains as the sole LIS plan. Once that plan is a monopolist of LIS enrollment, the LIS threshold equals the bid of that plan because it has captured all of the LIS auto-enrollees. The monopolist marks up its bid to the cost of the next closet competitor and earns a high profit because it is assigned all auto-enrollees.

There is a counteracting pricing incentive that prevents a complete unraveling. Plans that have a large share of lagged LIS enrollees have more inelastic residual demand curves above the LIS threshold. High lagged enrollment plans are more likely to price above the threshold than a plan with a similar cost position that has low or no lagged LIS enrollment. Those low lagged enrollment plans are more likely to price at the threshold. For the sake of completeness, the appendix contains a derivation of residual demand elasticities which makes this counteracting incentive apparent.

The combination of these two counteracting incentives induces a cycling effect, much like an invest-then-harvest pricing strategy, whereby plans go in and out of LIS eligibility year after year. Ericson (2014) shows how inertia in consumer choices of Part D plans induce invest-then-harvest pricing cycles. The mechanics of the LIS subsidy rules, even absent inertia in choice, contribute to cycling. Decarolis (2015) presents further reduced form evidence on these pricing incentives and considers further complications regarding the cycling effect.[13]

I present two pieces of basic evidence to support cycling. Table 3 presents time series serial correlations of premiums, basic/enhanced plan status, and LIS status from an AR(1) model estimated over the years 2006-2009. The second column reports the AR(1) correlation

[13]Decarolis (2015) documents a particularly dubious strategy in which insurers price their $t-1$ LIS plan above the LIS threshold to earn high markups while retaining control of the LIS threshold for that year. Concurrently, the insurer introduces a new, otherwise identical plan priced below the threshold to ensure control of the threshold in future years. Since 2012, CMS has instituted policies to prevent this practice.

Table 3: AR(1) Process of Product Characteristics

X	Constant	AR(1) coefficient	R^2
Monthly Premium	4.80*	0.95*	0.65
Enhanced Plan dummy	0.09*	0.90*	0.81
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 (2006-2009) and regions. *: estimates are significant at 1%

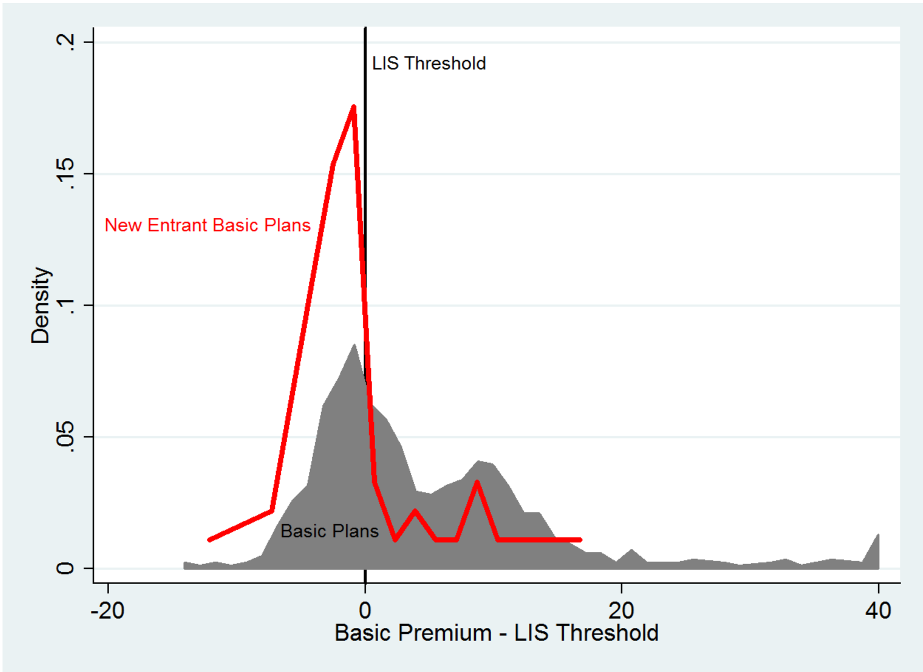

Figure 4: Bid Histogram: New Entrant Bunching @ LIS Threshold

between year $t-1$ and year t. For monthly premiums and basic/enhanced status there is a very high degree of correlation, greater than 0.90. The persistency in LIS status is much lower, 0.6, which is evidence of plans cycling in and out of LIS eligibility.

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 cycling into the market with LIS eligibility while the prior years LIS plans are cycling out of eligibility.

4.5 Consequences of Re-assignment on LIS Beneficiaries

The cycling of LIS eligibility can have adverse welfare consequences for auto-enrollees reassigned to new plans because it disrupts coverage continuity. LIS recipients face switching costs to conform to a new plan's formulary. They are a particularly vulnerable segment of the population given their very high usage of prescription drugs. Switching costs can be quite high. Miller and Yeo (2012); Polyakova (2015) estimate switching costs between \$1700 and \$2100 in Part D. Nosal (2012) estimates switching costs of \$4000 for Medicare Advantage plans. Theses estimates also include regular enrollees, who face switching costs attributable to the efforts of actively selecting plans. That sort of inattention bias does not contribute to the switching cost of re-assigned auto-enrollees. However, the estimates in the literature suggest the conforming cost of auto-enrollees could be quite high. The purpose of reassigning individuals out of above-threshold plans is to save on premium subsidy payments, but it may impose a very high welfare cost on re-assignees. Consumer advocates and CMS have given reassignment high priority for policy actions (Summer et al., 2010). The evidence shows that in most cases the subsidy dollar savings of reassignment are quite low, which prompted CMS to introduce the *DeMinimis* rule with the specific aim of reducing reassignment.

There may be health consequences. I analyze mortality, a leading health indicator, for 15,000 beneficiaries sampled from the administrative data. LIS beneficiaries who switch plans, primarily due to re-assignment, have a 23% higher mortality rate in the year following the switch than non-switchers, who were not re-assigned. For comparison, regular enrollee switchers, doing so by choice, have a 34% lower mortality rate than their peers who do not switch. These mortality statistics raise concern about a severe health consequences for re-assignees and, more generally, for those inattentive to plan selection. However, the statistics also suggest that health, in particularly impending mortality, contributes to choice frictions of plan selection.

5 Supply with Imperfect Risk Adjustments

The prior model of the supply-side abstracts away from risk selection by assuming a perfect risk adjustment mechanism. In this section, I expand the supply-side by explicitly modeling the three "R"s in Part D's risk adjustment mechanism. I present evidence that it is an imperfect risk adjustment model that gives insurers scope to cream skim relatively low risk enrollees. I then show how an imperfect risk adjustment mechanism is conjunction with the LIS rules leads to even more heavily distorted pricing.

5.1 Risk Adjustments and Risk Corridors

I follow Miller and Yeo (2013) to model risk adjustments. In the population there is heterogeneity in individuals' drug usage. Let person i's type be defined by the tuple $(\alpha_i, \beta_i, \kappa_i, r_i, a_i)$. The first three terms were defined in the demand model. The term $r_i > 0$ is an individual's risk score as assigned by Medicare. The risk score is the predicted value of drug expenditures as determined by Medicare's risk adjustment formula. The formula is estimated from a regression of drug expenditures on disease conditions (kidney failure, diabetes, etc.) and demographics (age, gender). It is a *prospective* risk scoring formula in that it is based on prior years data, and the formula is known by insurers at the time of submitting bids. Risk scores are normalized such that the average Medicare beneficiary enrolled in Part coverage has a risk score of 1. For example, someone with a risk score of 1.1 has a 10% greater than average predicted drug expenditure. The term a_i is called the *selection factor* which measures how an individual's actual drug expenditure deviates from that predicted by the risk scoring model.

Let c_{ij} denote the cost plan j incurs from enrolling a person of type i. Costs can be parameterized as,

$$c_{ij} = (c_j^{basic} + c_j^{enhanced})(r_i + a_i) \tag{20}$$

The parameters c_j^{basic} and $c_j^{enhanced}$ are plan specific costs factors that represent the plan's cost to cover an enrollee with average drug expenditures for the portions of coverage attributable to basic and enhanced cost sharing features. Costs scale in proportion to risk scores and selection factors. Enrollees with higher risk scores r_i are more costly. The selection factor, a_i, measures the cost difference of an individual from that predicted by his risk score. Individuals with positive values of a_i have higher costs than that predicted by the risk score; negative values, lower costs. I make a simplifying assumption that r_i and a_i have proportional effects on the basic and enhanced component of cost. This abstracts away from the non-linearity in cost sharing due to deductibles and the donut hole. A plan's *average* cost for its pool of enrollees can be determined by integrating across the distribution of types enrolled in the plan. Denote r_j and a_j as the average risk score and selection factor for plan j; thus $c_j^{avg} = (c_j^{basic} + c_j^{enhanced})(r_j + a_j)$. Note I am careful to not claim average cost equals marginal costs as was implicitly assumed in the supply-side model with perfect risk adjustments.

Medicare applies risk adjustments to adjust payments up or down based on plans' average risk scores, r_j and risk corridors to adjust payments up or down based on plans' average selection factors, a_j. The per enrollee average revenue R_j received by a plan that submits a bid of $(b_j^{basic}, b_j^{enhanced})$ with risk pool (r_j, a_j) is:

$$R_j = b_j^{basic} r_j + \theta c_j^{basic} a_j + b_j^{enhanced} + 0 c_j^{enhanced} a_j. \tag{21}$$

The first term is the risk adjustment payment, which scales the bid by the risk score. For example, the risk adjusted payment is 10% higher than the basic component of the bid for a plan with a risk score of 1.1. Payments are deducted from plans with risk scores less than 1. The second term is the risk corridor payment. Medicare deducts a portion θ of payments from plans that have selection factors less than 0, and compensates plans that have selection factors greater than zero. In other words, risk corridors act as a risk sharing scheme between the government and insurers to insure plans against the risk of enrolling a pool of individuals with realized costs deviating from that predicted by the risk scoring model. I assume a linear risk corridor parameter θ; the actual risk corridors use a step-wise function. Very few insurers are subject to no risk corridors, so the step-wise function becomes infra-marginal and the linearity assumption will have little bearing on the results. The third and fourth terms show that the enhanced component of the bid and cost do not factor into risk adjustments and risk corridors.

With linear risk corridors, the profit function for a plan with bid $(b_j^{basic}, b_j^{enhanced})$, risk score, r_j and selection factor a_j is

$$\pi_{jmt} = \left[(b_j^{basic} r_j + \theta c_j^{basic} a_j) - c_j^{basic}(r_j + a_j) + b_j^{enhanced} - c_j^{enhanced}(r_j + a_j)\right] s_{jmt}(\mathbf{b}) M_{mt}. \tag{22}$$

I define a perfect risk adjustment mechanism to be one in which $a_j = 0$ for all plans. Substituting $a_j = 0$ into the above equation, results in a profit function

$$\pi_{jmt} = \left[(b_j^{basic} r_j) - c_j^{basic}(r_j) + b_j^{enhanced} - c_j^{enhanced}(r_j)\right] s_{jmt}(\mathbf{b}) M_{mt}. \tag{23}$$

For a basic plan with $b_j^{enhanced} = c_j^{enhanced} = 0$, the profit equation is identical to the profit equation (10) presented in the perfect risk adjustment case; the r_j terms cancel out. This equivalence shows that a basic plan's pricing decision is not affected the composition of its enrollees' risk scores. Plans price as if they enroll an average cost beneficiary. For an enhanced plan, with $b_j^{enhanced} = c_j^{enhanced} > 0$, the equivalence does not hold; the r_j terms do not cancel out because the enhanced component of the bid is not risk adjusted.

A less stringent definition of a perfect risk adjustments does not require $a_j = 0$, but rather requires insurers' expectations of selection factors $E[a_j] = 0$ at the time of bidding. Even if the actual costs deviate, a risk neutral insurer would submit the same bid as if $a_j = 0$ with probability one.

Risk corridors affect bidding under two conditions. First, if risk adjustments are not perfect or improperly calibrated, selection factors could differ from zero. Insurers with positive selection factors price higher, and those with negative selection factors price lower. Second, if insurers are not risk neutral, they include a risk premia into their bids even if $E[a_j] = 0$. Medicare was concerned about both reasons at the outset of the program. They presumed that after three years when risk corridors were to phase-out, insurers should have sufficient experience that they behave as risk neutral and that imperfections in risk adjustments would stabilize. However, imbalances and volatility in risk adjustments persist through today, prompting Medicare to indefinitely continue risk corridors. In the next, I describe cream-skimming incentives and the connection to the LIS program that can explain the volatility.

5.2 Cream-Skimming

As can be seen in the profit equation (22), an insurer earns higher profit, all else equal, by enrolling a pool of low cost, low a_j enrollees. The incentive to attract low cost enrollees (or detract high cost enrollees) is called "cream-skimming." There is an important distinction between Part D and other non-risk adjusted insurance markets. Insurers want to cream skim with respect to selection factors, not risk scores. In non-risk adjusted insurance markets, plans want to cream skim with respect to the total cost of enrollees; that is, the sum of risk scores and selection factors $(r_j + a_j)$. For enhanced plans, profits are also higher if r_j is low because enhanced coverage is not risk adjusted, which give some incentive to cream skim with respect to risk scores.

The literature on risk adjustments in Part D describes how insurers adjust their benefit design to cream skim McAdams and Schwarz (2007); Hsu et al. (2009); Carey (2014) using formulary management techniques: drug exclusions, tiered copays, and usage restrictions. The risk scoring formula uses disease and demographics to predict drug expenditures. The regression models are not perfect predictors of drug expenditures; the R-squared values have ranged between 0.25 and 0.4 (Hsu et al., 2009). Cream-skimming techniques target specific drugs within therapeutic categories to attract the less costly patients of a particular disease and deter the more costly patients. For example the risk scoring model does not distinguish Type 1 and Type 2 Diabetics. A cream-skimming insurer would more favorably cover the drugs used by the less costly type with lower copay tiers, fewer exclusions, fewer restrictions, and unfavorably cover the more costly type. The regulations impose minimal restrictions on formulary construction, which gives insurers sufficient scope to differentiate coverage on a drug-by-drug basis. Provided enrollees compare formularies for their particular drug regimens an insurer can attract the more favorably selected patients. Carey (2014) also shows certain disease categories have become improperly scored over time due to changes in relative drug prices caused by patent expirations and new drug introductions. The risk score factors calibrations lag behind changes in pharmaceutical market prices between 4 and 5 years.

Figure 5 illustrates a stylized representation of how firms profit from cream skimming. The figure plots individuals' costs c_{ij} against risk scores r_i. The stars depict different people in the population with an unbiased regression line through the sample. Individuals above the regression line have $a_i > 0$; below the line, $a_i < 0$. The green oval circles the set of individuals a non-cream skimming insurer enrolls. In this example, the green insurer has $r_j = 1$; half of its enrollees are above, half below, the regression line. The red circle represents a cream-skimming insurer. Compared to the green insurer, it has a higher risk score r_j, yet has the same cost. The key point of the figure is to illustrate that the cream-skimming insurer has a greater share of its enrollees with $a_i < 0$, below the regression line. Both insurers have the same cost, but the cream-skimming insurer earns extra profit off of positive risk adjustment payments. Risk corridors would claw back some of this profit.

5.3 Risk Adjustment Imbalance and Volatility

Throughout the program's history, risk adjustment transfer payments have been out of balance and highly volatile. The actual costs incurred by insurers have not aligned with the cost predicted by the risk adjustment formula. In the aggregate, insurers' costs are significantly lower, implying that the risk adjustment formula over-estimates costs. At the disaggregated level, there are large differences across insurers. Some insurers have experienced deducted payments despite incurring high cost, while others have received payments despite having low cost. For any given insurer, there is a lot of volatility across years. Incurred costs relative to risk adjustments may be low one year, and then high the next year.

The risk adjustment formula underwent a major revision in 2011 that was intended to address the imbalances and volatility. From 2006 to 2010, the formula predicted costs based on drug expenditures of Medicare beneficiaries receiving drug coverage in the Federal Employee Health Benefit Program and Medicaid and extrapolated those predictions to the

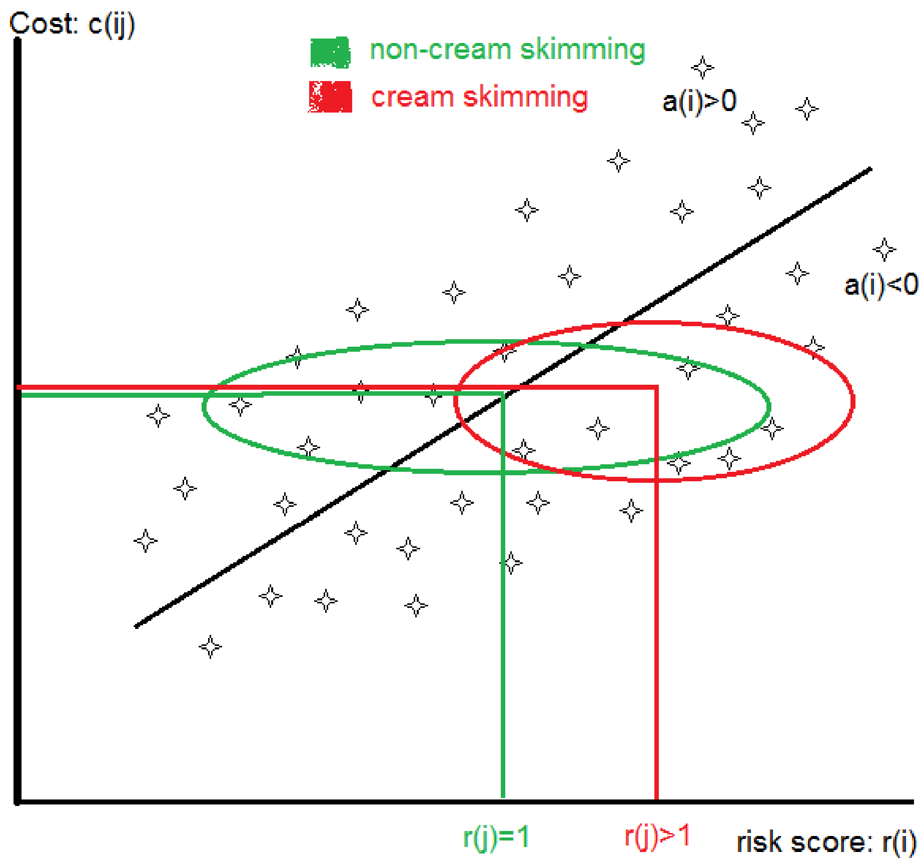

Figure 5: Cream skimming and non-cream-skimming insurer

entire Medicare population. This method was acknowledged as an imperfect method Hsu et al. (2010), but had to suffice because it was the best available data on prescription drug spending at the start of the program. The 2011 revision includes two major updates which are intended to better predict drug cost. First, predicted costs are based on historical drug claims data directly from Medicare Part D, not outside programs. The predictions are recalibrated every two or three years to keep up to date with changing market conditions. Second, the risk adjustment factors for each chronic condition are separated for LIS and non-LIS status because these populations have different drug spending habits, even for those with the same chronic conditions. There is also a separation for beneficiaries under age-65 because most of this population qualifies for Medicare through disability. Before 2011, LIS and under-65 beneficaries were grouped together with regular beneficiaries.

I use data on risk corridor payments to measure selection factors, a_j, and document the imbalance and volatility. Note from equation 21 that the risk corridor payment is $\theta c_j^{basic} a_j$. The data on risk corridor payments are aggregated to the insurer level, so it is not possible to calculate risk corridor payments accruing to any particular plan of a multi-plan insurer.

Table (4) reports statistics on risk corridor reconciliation payments aggregated across all plans for the years 2006-2012. The reconciliation process takes 2 years to finalize. Negative signs indicate risk corridors payments flowing from plans to CMS, which occurs if plans achieve favorable selection $a_j < 0$. The first column aggregates across both stand-alone Part D plans and Part D plans with bundled Medicare Advantage coverage. The second column aggregates across stand-alone Part D drug plans. Only stand-alone plans are eligible to receive LIS auto-enrollees. In 2006, the first year of the program, there was a tremendous amount of favorable selection. On a per-enrollee per month basis, stand alone plans paid out \$5.53 in risk corridor payments. Given an average bid of \$90 and $\theta \approx 0.8$, yields a conservative estimate for the selection factor of 7%. Based on the estimates of profit markups in table 2 ranging from 6% to 14%, favorable selection accounted for a very large proportion of profits. Similar calculations show Part D plans bundled with Medicare Advantage coverage profited even more.

After the first year, selection diminished and stabilized in 2008. However, in later years selection tilted negative again and increased in magnitude after the risk adjustment revision in 2011. The risk corridor parameter, θ, was higher in 2006 and 2007, so the implied degree of selection a_j in later years is comparable in magnitude to the 2006 figures.

Figure 6 depicts risk corridor reconciliation amounts between 2008 and 2012 for the drug plans (excluding bundled Medicare Advantage Part D) of seven of the largest national

Table 4: Aggregate National Risk Corridor Reconciliation Payments

year	All plans $ million	Stand-alone Part D $ million	Stand-alone Part D per person per month
2006	-$2,588	-$1,228	-$5.53
2007	-$599	-$204	-$0.95
2008	-$78	$100	$0.51
2009	-$795	-$500	-$2.01
2010	-$900	-$395	-$1.63
2011	-$902	-$412	-$1.84
2012	-$1,127	-$636	-$2.67
2013	-$737	*	*

2013 data only available in aggregate

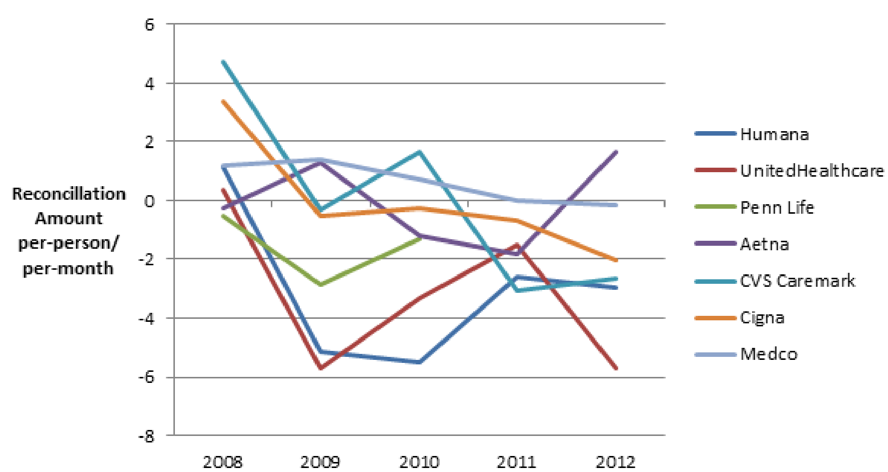

Figure 6: Risk Corridor Volatility

insurers. The amounts range from a low of \$-6 per-person per-month to a high of \$+5. There are two points to notice in this figure. In any given year, there are large differences across insurers. There is high volatility across years for any given insurer. The yearly ups and downs are generally not correlated movements for all insurers, but rather specific to an insurer. For example, between 2011 and 2012, Aetna and United Healthcare diverged, while Humana and CVS Caremark remained fairly stable.

The sustained negative risk corridors and volatility have been a puzzle for the Part D program. The revisions to the risk adjustment formula in 2011 generated some stability. The standard deviation in reconciliation payments across insurers reduced from \$5 per person per month in the 2008 to 2010 period to \$3 in 2011-2012. However, aggregate risk corridors payments remain large and negative. In correspondence to plan sponsors, CMS has attributed the patterns to a high degree of unpredictability in drug expenditures in the Part D program. In it's recent June 2015 report to Congress, MedPac has taken notice of the peculiar trends and is seeking explanations and policy recommendations.

5.4 Low Income Subsidies and Risk Selection

The puzzles about risk selection can be resolved by specifically considering the LIS population and rules regarding automatic enrollment. There are systematic patterns of selection linked to LIS enrollment. I first consider the risk sharing data and then model the effects of selection on pricing.

There is a systematic correlation between an insurer's LIS eligibility to receive auto-enrollees and its selection factor. Let $\theta_t c_{ft} a_{ft}$ denote risk corridor payments for all of the plans of insurer f in year t in per enrollee per month terms. I specify the following regressions in levels,

$$\theta_t c_{ft} a_{ft} = \beta(LISfrac_{ft}) + \beta(r_{ft}) + \alpha_f + \alpha_t + \epsilon_{ft}, \tag{24}$$

and in differences-in-differences,

$$\theta_t c_{ft} a_{ft} - \theta_{t-1} c_{ft-1} a_{ft-1} = \beta(LISfrac_{ft} - LISfrac_{ft-1}) + \beta(r_{ft} - r_{ft-1}) + \alpha_f + \alpha_t + \epsilon_{ft}. \tag{25}$$

The regressor $LISfrac_{ft}$ is the fraction of markets that insurer f operates in with an LIS eligible plan. The fractions are weighted by market size. α_f and α_t are insurer and year fixed effects. Insurer fixed effects control for differences across insurers in their ability to

Table 5: Risk Corridor Payments and LIS

	Levels				Differences-in-Differences			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LIS Fraction	2.52	2.68	3.21	3.59	5.98	5.37	7.55	6.51
	(1.11)	(1.21)	(1.26)	(1.42)	(1.95)	(1.59)	(2.57)	(2.01)
LIS Fraction (Post 2010)			-2.69	-3.78			-5.46	-4.57
			(1.26)	(1.56)			(3.02)	(2.57)
Risk Score		-7.80		-13.82		30.96		39.96
		(12.32)		(13.5)		(23.33)		(34.35)
Risk Score (Post 2010)				13.07				-28.80
				(8.49)				(38.37)
N obs	400	400	400	400	329	329	329	3.29

Standard errors in parentheses. Year and insurer fixed effects are included in all specifications.

favorably or unfavorably selection enrollees. Year fixed effects are important to control for the yearly trends in selection factors. The annual reconciliation payments in table 4 show large yearly swings in the aggregate. Year fixed effects also control for the reduction in risk sharing θ_t that occurred in 2008. The differences-in-differences specifications uses across time variation in whether an insurer is moving in or out of LIS eligibility. The cycling phenomenon discussed in the previous section generates a lot of across time, within insurer variation in LIS eligibility. I also include specifications that separate the effects for the periods before and after 2010 to test whether the 2011 revision to the risk adjustment mechanism altered selection patterns. The term r_{ft} is the insurer's risk score. In principle risk scores shouldn't matter for selection factors a_{ft} which are defined relative to risks scores. However there may be effects if the risk scoring model is a biased predictor of cost as suggested by Hsu et al. (2010), if higher risk score beneficiaries are more likely to be favorably selected as described by Brown et al. (2014) for Medicare Advantage, or if partial adjustments for enhanced plans lead to unfavorably selection. I drop all small insurers with fewer than 5000 enrollees, 88 total. This group has a very low rate of LIS coverage (11%) as compared to larger insurers (47%). There are large outliers in reconciliation payments which would be expected due to idiosyncratic insurance risks for small insurers.

Table 5 reports results for the specifications in levels and differences-in-differences. The coefficient on the LIS fraction is large and positive in all specifications. The positive sign indicates that insurers with many LIS plans have higher selection factors a_f than insurers with few LIS eligible plans. The magnitudes are larger, about double, in the differences-in-differences specifications that uses variation in insurers going in and out of LIS eligibility. These specifications best control for the tendency of insurers with a greater ability to favorably select to also be LIS eligible. In dollar terms, the range of magnitudes in the differences-in-differences specifications imply that an insurer going from no LIS eligible plans to all plans being LIS eligible, will experience between a \$5.37 to \$7.55 increase in risk corridor payments per enrollee per month. Using $\theta = 0.5$ as an approximate measure of risk sharing, the coefficient implies there is between a \$10.74 to \$15.10 increase in selection cost $c_{ft}a_{ft}$ when an insurer shifts between no LIS plans and 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 28% of all insurers). These selection differences account for a large proportion of costs and profits given the average bid is about \$90 and markups range from 6% to 14%.

The results about selection and LIS coverage are unaltered when controlling for risk scores, suggesting the patterns are not just an artifact of improper risk scoring of LIS beneficiaries. The specifications in levels show that higher risk score insurers have lower selection factors, implying higher risk score beneficiaries are more profitable. The magnitudes are small and statistically imprecise. A one standard deviation increase in risk scores of 0.1 lowers the reconciliation payments by \$0.78. The coefficient flips signs in differences-in-differences, suggesting that higher risk score beneficiaries are less profitable. Magnitudes are modest relative to those for LIS coverage. A one standard deviation increase in risk score differences of 0.04 raises reconciliation payments by \$1.24.

The revisions to the risk adjustment model in 2011 appear to have corrected some of the discrepancies in risk corridors. The interaction terms of LIS fraction and risk score for the post-2010 period indicate the systematic selection patterns for LIS coverage and risk scores found in the 2006 to 2010 period disappear after 2011. However, there are still large differences in risk corridor payments across insurers and the aggregate is highly negative.

The default plan assignment rules for LIS beneficiaries and cream-skimming may explain why LIS plans have higher selection factors than non-LIS plans. The formulary management techniques used by insurers to cream skim are only effective against enrollees who *actively* select plans. That is, they can cream skim regular enrollees. As illustrated in the demand estimates in section 3, LIS auto-enrollees are very unresponsive to plan features including the formulary management techniques used by insurers to cream skim. Furthermore, Carey (2014) shows that copays, more so than formulary exclusions, are the primary tool to cream-skim. The high copay subsidies for LIS beneficiaries eliminates copays as a cream-skimming tool against LIS beneficiaries. The random assignment rules distribute LIS enrollees uniformly across LIS eligible insurers. As a result, there would be a uniform distribution of a_i

types across LIS plans with mean $a_i = 0$, whereas cream skimmed regular enrollees would be distributed with $a_i < 0$. Similar ideas about risk selection of choosers and non-choosers have been studied in the context of switching costs (Handel, 2013; Polyakova, 2015). Enrollees locked-in by switching costs, like auto-enrollees, are more likely to opt for the default plan assignment which blunts selection.

The 2011 revision corrected the tendency for LIS beneficiaries to be unfavorably selected by creating separate risk adjustment models for LIS and non-LIS beneficiaries. The old formula included LIS and institutional status as a cost predictor to capture tendencies of these populations to have higher expenditures, independent of other disease and demographic factors. The calibration predicted 0.08 higher risk scores for LIS beneficiaries and 0.21 higher for institutionalized beneficiaries, many who are also LIS. Risk adjustment factors for each disease were the same regardless of LIS or institutional status. The new formula interacts LIS and institutional status with demographics and disease. Because cream-skimming occurs at the disease level, the new model better predicts cost for the regular and LIS beneficiaries. The data on risk scores indicates the new formula did not change relative risk scores between LIS and regular beneficiaries. Risk scores regressed on the fraction of LIS eligible plans for the pre-2011 period show that LIS eligible plans have 0.092 higher risk scores. The difference post-2010 is 0.098. I further explore the consequences of the revised risk scoring model in the section on reinsurance.

5.5 Pricing Distortions with LIS Subsidies and Selection

The supply-side response of insurers to LIS auto-enrollment changes when LIS enrollees are more costly on a risk adjusted basis than regular enrollees. Figure 7 modifies the prior figure to show the effect.

Now there is a discontinuity in marginal cost at the LIS threshold. Actively selecting enrollees, composing all of the enrollees for a plan priced above the LIS threshold, are low cost types (low a_i). Auto-enrollees, who are high cost types (high a_i), are assigned to plans priced at or below the LIS threshold. The higher cost of auto-enrollees causes marginal cost to abruptly increase at the threshold. Based on the estimates from equation 5, auto enrollees are \$7.55 more costly than actively selecting enrollees. The figure illustrates how two different insurers would price. The blue insurer has a low cost position (low c_j^{basic}). Despite the higher cost for auto-enrollees, the plan's cost position is sufficiently low for auto-enrollees to be profitable. The insurer would set its price just below the LIS threshold. The red insurer has a higher cost position. The jump in marginal cost for auto-enrollees puts the cost above

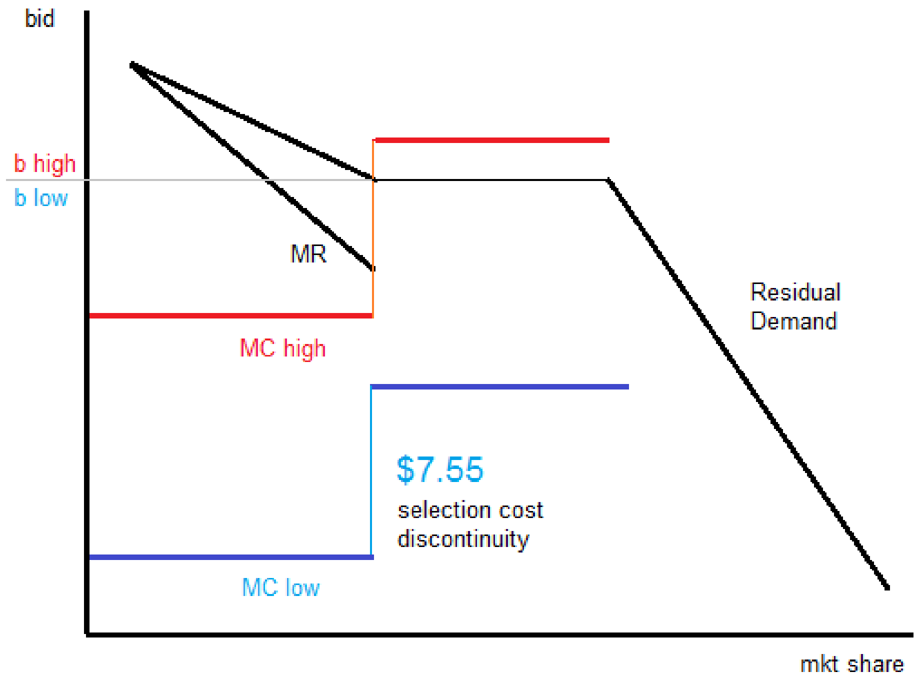

Figure 7: Pricing with Selection

the LIS threshold. This insurer wants to avoid auto-enrollees by pricing above the threshold. The figure has been drawn in a very specific manner. Notice the marginal revenue curve for the red insurer intersects the cost curve at the cost discontinuity. This plan optimally sets its price just above the LIS threshold. The demand and cost discontinuity implies that there are two groups of firms bunching prices at the threshold. Higher cost firms bunch right above, and lower cost firms bunch right below!

The pattern of bunching just above and just below is quite evident in kernel density plots of pricing. First, refer back to the density plots in figure 3. Careful inspection of the bids of basic plans shows a clear mass of plans pricing just above the threshold, within a couple dollars. The mode may be below the threshold, but the mass just above is still large. This is not necessarily complete evidence. The LIS threshold is not a known value to insurers at the time of bidding. If bidders have incomplete information about their rivals' cost and demand shocks, they cannot perfectly forecast the location of LIS threshold. The mass of plans pricing just above the threshold could be an artifact of imperfect information. Those plans may have intended to price below, but by chance accidently overbid.

Figure 8 depicts kernel density plots of basic plans for all years, 2007-2012, with the x-axis cropped to more easily view pricing around the LIS threshold. The blue plot shows plans that were LIS eligible in the prior (lagged) year; the red plots, plans that were not LIS eligible in the prior year. Given there is some persistency in costs, prior year LIS plans are more likely to be low cost plans in the current year than prior year non-LIS plans. New plans, like those in figure 4, are excluded. Prior year LIS plans have a distinct mode located \$1.60 below the LIS threshold. Imperfect information about the location of the LIS threshold causes the mode to be bounded away from the threshold. With perfect information, a plan that wants auto-enrollees would price exactly at the threshold, but, with imperfect information, such plans price with a safety margin a little bit lower. How much lower, depends on the degree of imperfect information. Although the mode is below the threshold, there is a large mass of prior year LIS plans pricing within 1 to 2 dollars above the threshold. The bounded mode below and mass of prices above when considering together help rule out imperfect information as the cause of plans having prices above. Apart from imperfect information, there are two explanations for why lagged LIS plans would price above. First, the theory predicts these plans may have been induced to price above because of a negative cost shocks. Second is the *Deminimis* rule which prevents lagged LIS plans from losing LIS auto-enrollees. However, the *demimimis* cannot be the sole explanation. There was actually a larger mass of plans pricing above in the years in which the *deminimis* rule was not in effect (2009,2010) than in the years it was in effect (2007,2008).

The pattern appears quite different for the plans that were not LIS in the prior year. There are two distinct modes around the LIS threshold. One is bounded \$1.60 below the threshold and the other is bounded \$1.00 above the threshold. Plans pricing below the threshold, profit off of auto-enrollees and have prices bounded away from the threshold because of imperfect information. Plans priced just above, want to set prices as low as possible (with a buffer), to avoid auto-enrollees. Between \$2 and \$5 above the threshold there is a dip in pricing, corresponding to the bid gap described in the theory illustrated in figure 2.

Strictly speaking, the theory illustrated in figure 7 predicts a mass of plans pricing exactly at the threshold and a non-zero mass of plans priced an infinitesimally small amount above. In a world with no or very little imperfect information, it would not be possible to statistically distinguish price bunching above from bunching below. The data would appear as one mass point. Incomplete information is the key to revealing bunching above and below. With incomplete information, both types of plans want to price close to the threshold, but not so close as to risk falling on the wrong side of the threshold. The theory predicts a "double hump" bounded away from the LIS threshold with the "dip" between humps occurs right at

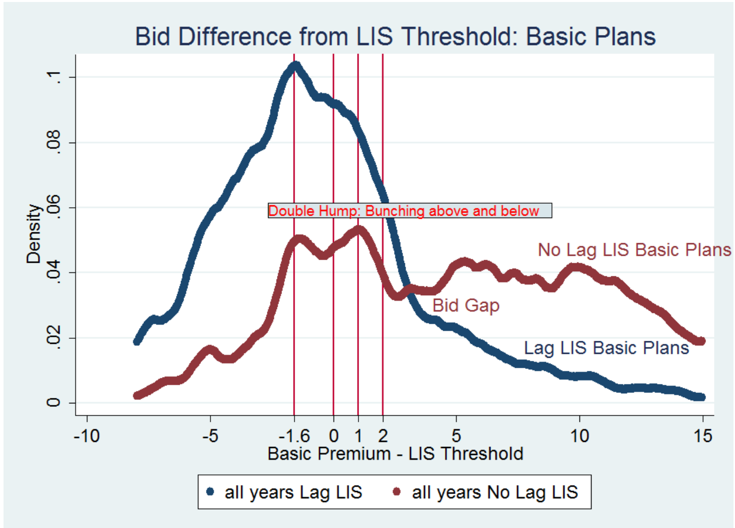

Figure 8: Bid Histogram: Bunching above and below LIS Threshold

the threshold. A Hartigan and Hartigan (1985) dip test of plans that were not LIS in the prior years (depicted in the figure), statistically rejects the hypothesis of a single mode (p-value 0.0327) in a local neighborhood of the LIS threshold.[14] Cropping the scale to a closer interval and reducing the kernel bandwidth, reveals "double humps" and "dips" located at the threshold across all years and for prior year LIS plans. The appendix includes more analysis. The deminimis rule complicates matters and generates a triple hump for prior year LIS plans.[15]

5.6 Does the LIS Threshold Rule Intensify or Soften Competition?: Revisited

The prior evidence from the model with perfect risk adjustments suggested softer competition because low cost insurers had an incentive to raise their bid up to the threshold. In a world with imperfect risk adjustments, there are plans priced just above the threshold. Here, the

[14]The dip test is a statistical test of multi-modality, that unlike the Silverman (1981) multi-modality test, does not rely on kernel densities. The dip test is a very conservative test against the null of a uniform distribution. Silverman's "bump hunting" test revealed the presence of many modes, up to 8. The high number of modes is likely a result of insurer and market heterogeneity in imperfect information regarding the location of the LIS threshold. Some markets or insurers may have more uncertainty, leading to modes farther away from the threshold.

[15]Prior year LIS plans have three choices. Price below the threshold to retain auto-enrollees and be assigned new auto-enrollees, price in the deminimis range to retain auto-enrollees, or price above the deminimis to shed auto-enrollees

competition effects are more transparent. The LIS threshold acts as a price floor. If it weren't for the LIS threshold, the plans bunching above would like to set lower prices. This can be seen in figure 7 by noting marginal revenue exceeds marginal cost at the cost discontinuity. The competitive effects are even worse considering that their higher prices act to increase the LIS threshold, thereby increasing the markups for low cost firms pricing just below the threshold.

The combination of a demand and cost discontinuity explain the teaser fact from the introduction about enhanced plans being priced lower than basic plans. There is a way for insurers to price lower and avoid LIS auto-enrollees: offer a (slightly) enhanced plan at a low price. Recall enhanced plans are ineligible for auto-enrollees no matter price. The high cost red insurer could offer an enhanced plan with the minimum allowed enhancement (eliminated deductible) at a price below the LIS threshold. Such a plan would have coverage characteristics "substantially similar" to basic plans and compete with other basic plans for actively choosing enrollees. Because all insurers must offer a basic plan, such an insurer would offer a "compliance" basic plan, priced above the LIS threshold. The low cost blue insurer with an LIS eligible plan can also follow this strategy. Note marginal revenue exceeds marginal cost at the threshold for both insurers in the example. However, the blue insurers basic plan is not a "compliance plan," rather it is an LIS plan. The pricing strategy becomes further complicated because some LIS enrollees choose plans and want to be avoided because of the risk adjustment imbalance. Suppose the LIS threshold is \$35. The insurer can set a very low price on the basic component (\$5) and a high price on the enhanced component (\$20), such that the price for a regular enrollee is \$30, lower than the \$35 price for LIS benchmark plans, but the subsidized LIS enrollee pays \$25, much higher than the \$0 price for LIS benchmark plans. This strategy explains the mass of enhanced plans illustrated in figure 3 with very low basic component prices and the peculiarity of enhanced plans having low cost in the structural cost estimates. There are pending proposals to eliminate this oddity of the market. In the policy recommendation section, I discuss each of these options and contend that it is better to address the selection discrepancies in the risk adjustment formula that give rise to these pricing strategies.

6 Reinsurance

The third "R" in the Part D risk sharing scheme is reinsurance. For enrollees with very high drug expenditures that reach "catastrophic" levels (\$5100 in 2006 and $\approx$ \$7062 in 2016), the

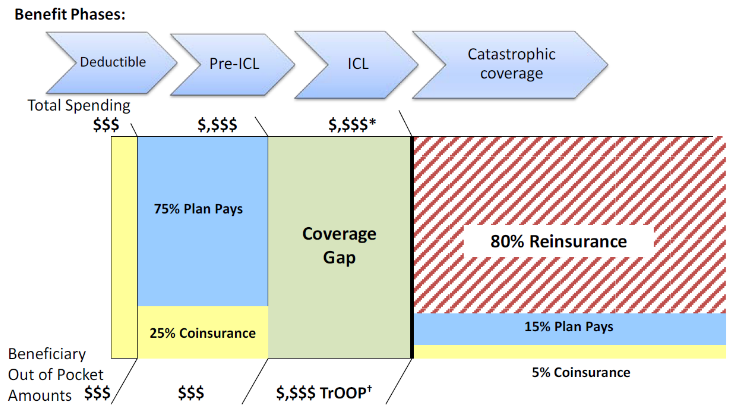

Figure 9: Sharing Risk with Government

government bears most of the cost. The enrollee pays a small copay, not to exceed 5%, the insurer pays 15%, and the government pays 80%. The portion paid by the government is called reinsurance.

Few regular enrollees reach the catastrophic level ($\approx$ 3%), many LIS beneficiaries do ($\approx$ 21%). LIS beneficiaries, in particular dual eligibles, have higher prescription spending because they have a higher incidence of disease, notable conditions include alzheimers (9$ regular, 19% dual eligibles), chronic kidney disease (14%, 20%), depression (11%, 27%) diabetes(25%, 36%) and heart failure (14%, 22%). On the upper tail, there are individuals with drug expenditures ranging into the hundreds of thousands of dollars.

Figure 9 depicts the share of drug expenditures paid by the insurer, enrollee, and government as a function of a patient's drug expenditures for a defined standard basic plan. Several "attachment" points—deductible, Pre-icl (pre-donut hole), donut hole, catastrophic—demarcate threshold points where shares change. I focus on reinsurance and, for the sake of this paper, do not model the nuances of the non-linearity at earlier attachment points.

With reinsurance factored in, it is necessary to re-model the cost equation. The concept of reinsurance is to place a cap c_j^{cap} on the cost incurred by the insurer. The cost of enrolling an individual of type (r_i, a_i) becomes

$$c_{ij} = min\{(c_j^{basic} + c_j^{enhanced})(r_i + a_i), c_j^{cap} + 0.15(c_j^{basic} + c_j^{enhanced})(r_i + a_i - (r+a)_j^{reinsurance}\} \tag{26}$$

The first element in the min function represents the cost born by the insurer for a low cost (low $r_i + a_i$) individual. The second element represents cost for types that have values of

$(r_i + a_i)$ exceeding the catastrophic threshold. The insurer pays its full share of cost up to the cap c_j^{cap} plus a 15% share of the cost exceeding the cap. The value $(r+a)_j^{reinsurance}$ is defined to be the range of $r_i + a_i$ for plan j that triggers catastrophic spending. Average drug expenditures are about half the catastrophic level, so $(r+a)_j^{reinsurance} \approx 2$. The main feature of reinsurance that I emphasize is the cap. To simplify exposition I will refer to this equation under the assumption that the insurer share of cost above the cap is zero and consider just basic plans ($c_j^{enhanced} = 0$).

With reinsurance, the *actual* per enrollee cost c_j^{actual}, is no longer the simple average of r_i and a_i. Instead cost is bounded by the cap. Determining the actual cost requires integrating over a truncation of the joint distribution of $F_j(r_i, a_i)$ of types enrolled in plan j. For a basic plan

$$c_j^{actual} = \int (c_j^{basic})(r_i + a_i)dF(r_i, a_i) - \int_R (c_j^{basic})(r_i + a_i)dF(r_i, a_i | r_i + a_i > (r+a)_j^{reinsurance}). \tag{27}$$

The first term is the simple average across risk scores and selection factors. The second term is the truncation for reinsurance where the region of integration R is defined over the range of (r_i, a_i) in the support of the distribution of F such that $(r_i + a_i) > (r+a)_j^{reinsurance}$.

The *predicted* per enrollee cost $c_j^{predicted}$ at the time of bidding is an actuarial calculation indexed to Medicare's risk adjustment formula.

$$c_j^{predicted} = \int (c_j^{basic})(r_i)dF(r_i) - \int_R (c_j^{basic})(r_i)dF(r_i | r_i > r_j^{reinsurance}). \tag{28}$$

The difference between the actual and predicted calculation is that the predicted amount assumes no selection ($E[a_i] = 0$). The prior theory and results about cream-skimming and unfavorable selection of LIS enrollees shows that predicted and actual costs can differ. This implies that predicted and actual reinsurance payments (the truncations) will differ. Medicare makes prospective reinsurance payments to plans based on predicted costs, which are then reconciled two years later based on actual cost.

For $\kappa_i = 1$ individuals the government bears most of the enrollee's share of drug costs. Plans are paid by the government the share of LIS copay subsidy costs, called Lics (low income cost sharing), in the deductible, initial coverage, donut hole, and reinsurance regions. Figure 9 shows that a large proportion of the lics payments occurs in the donut hole region because the plan bears zero cost. A high proportion, 42%, of LIS beneficiaries reach donut hole spending levels, whereas only 20% of regular enrollees reach those levels. Prospectively payments for Lics are made based on actuarial calculations that integrate across the spending distribution under the assumption of no selection ($E[a_i] = 0$). The integration formula is quite complicated because of the non-linearity in cost sharing, but follows the same principle as the integration for reinsurance. Lics costs are reconciled 2 years later based on actual cost. If there is selection, prospective and actual Lics payments will differ. There are no prospective payments for risk corridors because the predicted risk corridor payment is zero by definition when calculations are based on the assumption that $E[a_i] = 0$.

6.1 Reinsurance: A Growing Cost Burden

Table 6 reports aggregate reconciliation payments for risk corridors, Lics, and reinsurance. The payments represent the difference between the actual and predicted reinsurance costs, $Reins^{rec}_{jt} = Reins^{actual}_{jt} - Reins^{predicted}_{jt}$, and Lics costs as described by equation (27) and (28). Like risk corridor payments, reconciliation payments for reinsurance and Lics have exhibited wide variation throughout the program's history. In 2006, reinsurance reconciliation was large and negative ($1.54 billion) indicating that actual reinsurance costs were much lower than predict reinsurance costs. By 2008, reinsurance reversed, actual costs were $1.2 billion higher than predicted costs. Reinsurance stabilized in 2009. Lics payments also experienced a large spike of $1.2 billion in 2008, then reduced.

A new trend has emerged after the risk adjustment formula revision in 2011. Between 2010 and 2013, reinsurance reconciliation payments have increased nearly 10 fold from $549 million to $4.9 billion. Lics reconciliation increased 5 five fold from $332 million to $1.55 billion. The trend shows that actual costs chronically exceed predicted costs. The gap is growing rapidly. The trend for risk corridors is the opposite: actual costs consistently come in lower than predicted costs. This is quite puzzling. One measure of cost is below expectation, while the other is above.

The last two columns of table 6 report predicted reinsurance cost and actual reinsurance cost. Predicted reinsurance steadily grew a total of 60% from 2007 to 2012, outpacing enrollment growth of 30%. The trend sharply increased afterwards. Between 2012 and 2015, predicted reinsurance grew 225%, while enrollment grew 24%. At issue is not just the growing discrepancy between between actual and predicted reinsurance, but also the total growth in reinsurance.

There is another new trend in the market parallelling the rise in reinsurance. Bids (modeled as $\bar{b}_{jt}$ in equation (13)) are rapidly falling. Table 7 reports per enrollee per month dollar figure trends for reconciliation payments, reinsurance, premiums, and bids. Monthly

Table 6: Reconciliation Payments: Risk Corridors, Low Income Cost Sharing, Reinsurance

	Reconciliation Payments			Reinsurance	
	Risk Corridors	Lics	Reinsurance	Predicted	Actual
2006	-2,588	78	-1,544	8,165	6,621
2007	-599	406	186	8,309	8,495
2008	-78	1,241	1,206	9,507	10,713
2009	-795	376	-66	12,005	11,939
2010	-900	332	549	13,111	13,660
2011	-902	342	1,547	14,888	16,435
2012	-1,127	633	3,182	13,237	16,419
2013	-737	1,559	4,915	19,341	24,256
2014	*	*	*	24,232	*
2015	*	*	*	29,798	*

2013 data only available in aggregate. All values in millions of dollars.

Table 7: Payments Per Person Per Month

	Reconciliation Payments			Reinsurance		Premiums, Bids			Bids+Reins	
	Risk Cor.	Lics	Reins	Predicted	Actual	$\bar{p}$	$\bar{s}^g$	$\bar{b}$	Predicted	Actual
2006	-9.11	0.27	-5.44	29	23	32	60	92	121	115
2007	-1.94	1.32	0.60	27	28	27	53	80	107	108
2008	-0.24	3.78	3.67	29	33	28	53	81	110	114
2009	-2.32	1.10	-0.19	35	35	30	54	84	119	119
2010	-2.54	0.94	1.55	37	39	32	56	88	125	127
2011	-2.42	0.92	4.16	40	44	32	55	87	127	131
2012	-2.81	1.58	7.93	33	41	31	53	84	117	125
2013	-1.64	3.47	10.93	43	54	31	48	79	122	132
2014	*	*	*	51	*	32	43	75	126	*
2015	*	*	*	60	*	33	37	70	130	*

All values in dollars per person per month.

bids were high 2006 ($92), then reduced to $80 in 2007 before starting a gradual rise to a peak of $88 in 2010. Bids then started a decline to $70 in 2015. Throughout this time period premiums remained steady in the low thirties. A premium decline did not accompany the bid decline because the subsidy ratio λ_t declined. The subsidy ratio is calculated as 0.745 less the predicted reinsurance cost. By 2015 the subsidy ratio fell to 0.52 because of the rise in predicted reinsurance. Overall program cost, factoring in both the bid and reinsurance, have grown from a trough in 2007 of $108 to its peak of $130 for 2015 predicted cost, and $132 for actual cost in 2013, the most recent year that has completed the reconciliation process. These trends present conflicting patterns on cost. On one hand, premiums have remained stable and insurers are lowering their cost. On the other hand, reinsurance is rapidly growing and driving program costs higher.

Table 8: Reinsurance Reconciliation Payments and LIS

	Levels				Differences-in-Differences			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LIS Fraction	7.65	6.18	8.06	6.55	8.47	5.81	8.85	4.09
	(2.74)	(2.54)	(3.07)	(2.88)	(2.31)	(2.53)	(2.98)	(3.64)
LIS Fraction (Post 2010)			-1.60	-1.45			-1.31	2.68
			(3.48)	(3.30)			(5.29)	(6.64)
Risk Score		70.89		80.17		133.78		185.75
		(27.81)		(31.96)		(51.56)		(78.56)
Risk Score (Post 2010)				-13.90				-119.55
				(18.83)				(147.23)
N obs	400	400	400	400	329	329	329	329

Standard errors in parentheses. Year and insurer fixed effects are included in all specifications.

6.2 LIS and Reinsurance

In this section I seek to reconcile the patterns of reinsurance and program cost trends by focusing on the LIS program. I first present evidence from the reconciliation data showing how reinsurance is linked to LIS plan assignment. I then turn to micro-data on prescription drug prices from administrative data.

I apply the same regression model to reinsurance reconciliation payments used in equation 24 and 25 for risk corridor payments to test whether LIS auto-enrollees are unfavorably selected (higher a_i) relative to regular enrollees. I also separate the effects for the post 2010 period to test whether the risk adjustment model re-aligned payments. Table 8 reports the results. In all specifications there is a positive relationship between reinsurance and LIS coverage as was found between risk corridors and LIS coverage. A change from being a non-LIS plan to an LIS plan increases reinsurance cost in a range from \$5.81 to \$8.85 relative to that predicted by the risk scoring model. These dollar figures are on the same scale as that estimated for risk corridors. Higher risk scores also predict higher reinsurance, but only slightly reduce the effect attributable to LIS coverage. In the Post 2010 period, after the risk adjustment revision, the positive relationship between LIS coverage and reinsurance persists. The results on risk corridors showed that the relationship disappeared after the revision. This evidence suggests the revision did not effectively address the discrepancies in the reinsurance range of spending.

I estimate another specification that controls for reconciliation payments of risk corridors and lics reconciliation payments:

Table 9: Reinsurance Reconciliation Payments: Additional Controls

	Levels			
	(1)	(2)	(3)	(4)
LIS Fraction	6.18	5.60	6.55	5.18
	(2.54)	(2.67)	(2.88)	(2.98)
LIS Fraction (Post 2010)			1.45	1.69
			(3.30)	(3.47)
Risk Score	70.9	60.6	80.2	78.0
	(27.8)	(19.0)	(32.0)	(21.5)
Risk Score (Post 2010)			-13.9	-30.8
			(18.8)	(15.4)
Lics		0.62		0.62
		(0.12)		(0.13)
Lics (Post 2010)				0.01
				(0.20)
Risk Corridor		-0.06		-0.07
		(0.20)		(0.22)
Risk Corridor (Post 2010)				0.44
				(0.37)
N obs	400	400	400	400

Standard errors in parentheses. Year and insurer fixed effects are included in all specifications.

$$\begin{aligned} Reins^{rec}_{ft} &= \beta(LISfrac_{ft}) \\ &+ \beta(Lics^{rec}_{ft}) + \beta(RiskCor^{rec}_{ft}) + \beta(r_{ft}) + \alpha_t + \epsilon_{ft}. \end{aligned} \tag{29}$$

In this specification, any deviations in selection factors should be fully captured by the additional controls. That is, if an insurer's selection factor a_f increases because LIS beneficiaries are unfavorably selected, then $RiskCor^{rec}_{ft}$ and $Lics^{rec}_{ft}$ would increase. Conditional on those terms, the fraction of LIS plans should have zero effect on reinsurance reconciliation.

Table 9 compares estimates. The additional controls have negligible effect on the result for LIS coverage. This implies that the upper tail of LIS enrollees reaching reinsurance spending are more costly even after controlling for actuarial discrepancies in risk scoring and selection factor differences due to cream-skimming, . Aggregated to the national level for 2012, these estimates point towards a $2.8 billion gap in unaccounted for reinsurance payments.

6.3 Accounting for Excess Reinsurance Payments: Drug Price Discrimination

There is an explanation for the excessive reinsurance payments to LIS plans which goes to the heart of problems with government risk sharing programs. Once a patient enters catastrophic expenditures, the government pays most of the price for that patient's drugs. Yet, the government does not set the price of those drugs. Insurers and drug manufacturers in the Part D market have free reign to set prices with minimal government interference. Negotiations between these private parties determine prices.

Duggan and Scott-Morton (2010) show that Part D's system of relatively unencumbered negotiations has, for the most part, resulted in lower drug prices as compared to retail despite Part D being a highly subsidized marketplace. Insurers have a strong incentive to keep drug prices low when either the insurer or its enrollees, who select plans based on premiums and coverage characteristics, bear the cost. The flexibility in formulary management gives insurers the bargaining tools necessary to keep prices low. This general success of Part D is prone to fail when it comes to reinsurance and the LIS program. Drug manufacturers have an incentive to raise drug prices and insurers have little incentive to keep them low when the government bears the cost. Particularly high prices for LIS plans and drug claims covered under reinsurance could explain the findings of excessive reinsurance payments at the aggregate level.

At the micro-level, I analyze drug prices from CMS prescription drug event (PDE) claims records to test whether drug prices are higher for LIS beneficiaries and for claims paid under reinsurance. Every prescription filled at the pharmacy generates a claims record with detailed information about the type of drug, price, insurer, patient, prescriber, pharmacy, and cost sharing. I access PDE records through CMS's Chronic Conditions Data Warehouse (CCW) for a 5% sample of the Medicare population. The database contains hundreds of millions of claims.

To test the hypothesis about drug prices, I regress log drug prices $\log(p_i)$ on indicators variables for whether the event, i, is the claim for an LIS beneficiary LIS_i and a set of indicator variables for which part of the benefit phase depicted in figure 9 the claim is being paid under: deductible, initial coverage zone, donut hole, reinsurance. The excluded category is the initial coverage zone, which is the region in which the insurer bears the highest share of cost (75%). Price is measured as the price per-day of prescribed treatment. Higher prices for claims covering LIS beneficiaries or reinsurance events support the hypothesis of drug price discrimination targeting claims in which the government bears a large share of cost.

$$\log(p_i) = \beta LIS_i + \beta(deductibleevent) + \beta(donutevent) + \beta(reinsuranceevent) + f + \epsilon_i \quad (30)$$

I consider several specifications that include fixed effects f for the type of drug, insurer, pharmacy where the prescription is filled, patient identifiers, and calendar dates. In specifications without calendar date fixed effects, I include a linear time trend $\beta * date$ because prices for any particular drug tend to depreciate over time. The PDE records report actual transaction prices that are net of all rebates negotiated by the insurer to be applied directly to the claim. Insurers may also receive manufacturer rebates that are not applied to the claim. I further discuss these rebates in the next section. Because transaction prices are proprietary, I am not authorized to link PDE records with any other data so as to protect the commercial integrity of the Part D program. Patient and insurer identities are encrypted for privacy. I infer from cost sharing information whether the patient is an LIS beneficiary by noting whether the government pays part of the claim (lics amount). I categorize a patient as LIS if he has at least one claim in a calendar year with a lics payment. All reported results have been approved for dissemination by CMS. I perform estimation on all claims for a randomly selected sample of 15,000 beneficiaries with at least 1 claim in 2010 over the years 2008 through 2011. The resulting sample size is 2,054,051 prescription drug events. [16]

Tables 10 11 and 12 report results. The first set of results in table 10 excludes fixed effects for drug type. The remaining specifications include more fine-grained fixed effect controls for the type of drug, so that price differences are identified off of within drug type variation in prices. Excluding fixed effects, the first specification shows that drug prices for LIS beneficiaries are 0.34 log points (40%) higher than for regular beneficiaries. This could be evidence of price discrimination targeting higher prices for LIS beneficiaries, but it could also indicate that LIS beneficiaries take the varieties of drugs that are more expensive. The next specification includes a control for the patient's total drug expenditures measured as the annual amount drug spending on other prescription fills. Controlling for total drug expenditures, LIS beneficiaries face prices 6.4% higher than regular enrollees. I include total spending controls in many other specifications. The 3rd and 4th specifications consider drug

[16] I sample 15,000 beneficiaries with at least 1 prescription drug event in the 2010 calendar year. I include events for those same people in the years 2008 through 2011. This sample size is largest that can be estimated given the computing capabilities of the CCW workstations. It is sufficiently large for statistical purposes. I exclude all non-Part D drugs (i.e. prescription sleep aids) separately covered as enhancements and prescriptions compounded with more than 1 drug ingredient. These exclusions account for about 1% of claims.

Table 10: Drug Price Discrimination

	No Drug Fixed Effects				Drug Ingredient Fixed Effects			
LIS Beneficiary	.340	.064	.219	.068	.093	.059	.058	.037
	(.002)	(.002)	(.002)	(.002)	(.007)	(.007)	(.004)	(.003)
Deductible Event			-.411	-.165	-.093	-.039	-.068	-.034
			(.004)	(.004)	(.007)	(.006)	(.005)	(.004)
Donut Event			.393	.120	.088	.026	.052	.016
			(.003)	(.003)	(.005)	(.004)	(.003)	(.002)
Reinsurance Event			.747	.261	.154	.047	.088	.023
			(.004)	(.004)	(.014)	(.012)	(.005)	(.005)
Log Other RX		.332		.279		.067		.042
Annual Spend		(.001)		(.001)		(.004)		(.002)
Date Trend (x1000)	-.158	-.156	-.220	-.178	-.169	-.159	-.144	-.139
	(.002)	(.002)	(.002)	(.002)	(.034)	(.034)	(.023)	(.023)
				Fixed Effect Categories				
DaySupply-Quantity							Y	Y
-Dosage Form							Y	Y
N. F.E. Categories					1228	1228	51355	51355
N. Drug Events	2,054,051	2,054,051	2,054,019	2,054,019	2,054,019	2,054,019	2,054,019	2,054,019

prices in the different phases of the benefit. The excluded category is the initial coverage zone when the insurers pays 75% of the claim. Reinsurance and donut hole claims are much more costly (111% and 48%); LIS claims remain high (25%), while deductible claims are less costly (-34%). The scale of these effects drops by a factor of 4, but remains high when controlling for the patient's total drug expenditures. In this and all other specifications that include fixed effects, prices tend to be lower in the deductible region, higher in the donut hole region, and the highest under reinsurance.

The last four columns include fixed effect for the active drug ingredient. Price effects are identified off of price variation within a type of drug, not across very diverse set of drugs. The pattern persists that LIS beneficiaries and events covered under reinsurance have higher prices. Prices are higher by about 10% for LIS beneficiaries and 15% for reinsurance claims. The last two columns include more detailed fixed effects for the days-supply and dosage form. Typically, longer days supply (90day vs 30day) gives some price discount on the per-day price. Different dosage forms, (tablet, capsule, extended release, etc.), for the same drug ingredient may have different prices. The LIS and reinsurance effects are smaller, but still show prices differences of 5.8% and 8.8% when the comparison is for the same drug ingredient of identical day-supply and dosage form. If there is some degree of clinical substitutability between prescriptions with long and short days supply and among dosage forms, the attenuation in the price effects reveals how drug manufacturers are able to price discriminate against types of claims. LIS beneficiaries and those reaching reinsurance are more likely to take expensive dosage forms and fill for shorter days-supply. But if there is little clinical substitutability, it would not necessarily be price discrimination, rather different days supply and dosage forms represent unique treatment regimens.

The next set of results in table 11 control for all conceivable differences in the pharmacological characteristics of the drug: drug ingredient, strength (5mg vs 10mg), dosage form, Branded/non-branded. The pharmacy consulting company First DataBank provides the classifications for CCW. This information is only available for 2010 and 2011, so the sample size falls by about a half. The first column shows the price differences still appear for LIS and reinsurance, 9.7% and 9.1%. The next columns explore whether the price differences can be explained by differences in the insurer administering the claim. Not all insurers can bargain for the same prices, and it may be that insurers with more LIS or reinsurance claims have less ability to negotiate low prices. With insurer fixed effects, the price differences remain. That is, the LIS beneficiaries of a particular insurer face higher prices than that insurer's non-LIS beneficiaries. Reinsurance claims for that insurer are also higher. The third column includes the plan in the fixed effect because insurers offer many plans. Different plans of the same insurer can have different formularies, cost sharing, and usage restrictions. As the theory about LIS bunching shows, they can also have different proportions of LIS and non-LIS beneficiaries. With insurer-plan controls the result persists. Claims for patients in the exact same plan (in the same region) are higher for LIS and reinsurance prescription fills. The fourth column has a pharmacy outlet fixed effect. Note that unique retail locations of large chain pharmacies are considered different pharmacies. The effects are still positive, but smaller. The attenuation suggest a major channel of price discrimination is that pharmacies serving a greater share of LIS beneficiaries and patients reaching reinsurance price higher. Pharmacies serving institutionalized patients are a likely source of the higher prices. The fifth column includes fixed effects for a particular individual in a year. An individual is enrolled in the same plan for the entire year, except in exceptional circumstances. Consumers usually fill prescriptions from the same pharmacy. This is the smallest reinsurance effect, 2%, but suggests that something closer to 1st degree price discrimination against a beneficiary occurs in the market. The final column is the most heavily controlled and perhaps most interesting. It includes a control for exact calendar dates. Reinsurance events follow the other events (deductible, initial coverage, donut hole) in calendar time. There are more reinsurance events in December than in January. There may be some seasonality in drug price explaining the price differences. The result shows that an LIS beneficiary with the same insurance company, filling the exact same prescription, on the same calendar day has 4% higher drug price. It is 3% higher for reinsurance claim.

Table 12 is similar, except the drug type fixed effect is based on (NDC) national drug codes. The NDC is unique in the pharmacological attributes, but differs across "labelers", who are the manufacturers or distributors of the drug. There are examples of drugs, even branded drugs, that have many manufacturers and distributors. The results are very similar

Table 11: Drug Price Discrimination

	Formulary Code Fixed Effects (Ingredient/Strength/Dosage Form/Brand Unique)							
LIS Beneficiary	.096	.082	.098	.010		.054	.043	.040
	(.006)	(.004)	(.007)	(.014)		(.004)	(.003)	(.012)
Deductible Event	-.065	-.063	-.035	-.004	-.004	-.050	-.039	-.058
	(.005)	(.003)	(.003)	(.002)	(.002)	(.004)	(.003)	(.023)
Donut Event	.068	.051	.035	.009	.008	.041	.025	.034
	(.004)	(.002)	(.002)	(.002)	(.002)	(.003)	(.002)	(.016)
Reinsurance Event	.091	.076	.058	.024	.020	.053	.038	.033
	(.007)	(.004)	(.004)	(.003)	(.003)	(.005)	(.003)	(.022)
Date Trend (x1000)	-.147	-.206	-.162	-.088	-.082	-.115	-.128	
	(.030)	(.011)	(.008)	(.006)	(.006)	(.023)	(.009)	
	Fixed Effect Categories							
DaySupply-Quantity						Y	Y	Y
Insurer-Year		Y	Y				Y	Y
-Plan			Y					
Pharmacy-Year				Y				
Person-Year					Y			
Calendar Date								Y
N. F.E. Categories	2,881	101,172	197,446	279,316	261,011	45,224	210,796	1,007,964
N. Drug Events	1,059,120	1,059,120	1,059,120	1,058,991	1,059,120	1,059,120	1,059,120	1,059,120

Table 12: Drug Price Discrimination

	National Drug Code (NDC) Fixed Effects (Ingredient/Strength/Dosage Form/Manufacturer Unique)							
LIS Beneficiary	.070	.045	.043	.012		.040	.025	
	(.003)	(.003)	(.006)	(.011)		(.002)	(.002)	
Deductible Event	-.057	-.035	-.014	-.004	-.002	-.044	-.015	-.004
	(.003)	(.002)	(.002)	(.001)	(.001)	(.002)	(.001)	(.001)
Donut Event	.059	.036	.019	.012	.009	.034	.014	.007
	(.002)	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)	(.001)
Reinsurance Event	.089	.065	.038	.028	.023	.049	.026	.018
	(.004)	(.003)	(.002)	(.003)	(.002)	(.002)	(.002)	(.001)
Date Trend (x1000)	-.077	-.136	-.086	-.068	-.056	-.053	-.063	-.046
	(.012)	(.006)	(.004)	(.005)	(.004)	(.010)	(.004)	(.003)
	Fixed Effect Categories							
DaySupply-Quantity						Y	Y	Y
Insurer-Year		Y	Y				Y	Y
-Plan			Y					
Pharmacy-Year				Y				
Person-Year					Y			Y
Calendar Date								
N. F.E. Categories	14,894	40,117	535,470	603,650	592,818	132,127	546,987	703,949
N. Drug Events	2,054,051	2,054,051	2,054,051	2,053,643	2,054,051	2,054,051	2,054,051	2,054,051

to those that don't have the labeler distinction.

In summary, these results reveal pattern of high prices for reinsurance events and LIS beneficiaries. The results match the prediction that they should be higher because the government bears in the cost, yet doesn't set prices. The results also show drug suppliers vis-a-vis insurers are able to price discriminate across claims. Price discrimination occurs along several dimensions. Drugs with the same active ingredient can be prescribed in a variety of dosage forms, strengths, quantities of pills, days-supply, and branded/non-branded varieties. If all of these combinations are truly unique drugs, with no clinical substitutability, then the results should not be considered price discrimination, rather different prices for different products. But, if there is some substitutability, the results indicate a price discrimination scheme targeting high prices for the varieties filled by LIS beneficiaries and patients reaching reinsurance spending levels. Price discrimination also occurs across insurers, within insurers across their different plans, across pharmacies, and across calendar dates. I interpret this as third degree price discrimination, where the industry participants have identified differences in patients in the plans they enroll in and places and times at which they fill prescriptions. The auto-assignment rules aid in price discriminating by separating LIS and regular beneficiaries according to whether the plan prices above or below the threshold. Finally, there may be evidence of first degree price discrimination targeting specific claims of individuals.

6.4 Will Drug Price Discrimination Unravel the Part D Market

Price discrimination against reinsurance claims and LIS beneficiaries may explain the excessively high reinsurance for insurers with many LIS eligible plans. But there may be additional factors contributing to the more general trend of rising reinsurance and declining bids. In this section, I discuss how manufacturer rebates, the 2011 risk adjustment revisions, and minimum loss ratio requirements that took effect in 2014 may be interacting with drug price discrimination to accelerate the trends and potentially unravel the market.

Manufacturer rebates may be a critical element necessary to reconcile the trends. Insurers negotiate with pharmacies, wholesalers, and manufacturers over the drug prices charged to beneficiaries at the point of sale, the values reported in PDE records. Drug manufacturers and other suppliers offer rebates to insurers that are not applied at the transaction. A 2008 DHHS Office of the Inspector General reported \$6.5 billion in 2008.[17] Insurers submit quarterly and annual rebate reports to CMS an aggregated level. The rebates decrease the insurer's share of cost (modeled as c_j) and reinsurance. Enrollees do not receive a rebate check for their share of cost. Rebates factor into reconciliation payment data but do not appear in PDE records made available to researchers.

The transaction level data likely understates the magnitude of price discrimination. Insurers and drug suppliers can strategically apply rebates to maximize the surplus extracted from marking up drug prices on reinsurance and LIS claims. The surplus maximizing rebate scheme would apply rebate dollars to claims in which the insurer and regular beneficiaries bear cost and minimize rebates applied to claims where reinsurance applies. For example, to exploit seasonality, rebates could be offered early in the calendar year when few beneficiaries have reached catastrophic spending to prevent rebate dollars from being applied to reinsurance. To exploit the segmentation of regular and LIS enrollees, a multi-plan insurer can apply rebate dollars to its non-LIS plans. Many other more sophisticated rebate scheme could be devised. For a manufacturer offering multiple drug types, rebates could be applied for drug disproportionately used by regular beneficiaries or low spenders. Insurers can use

[17] Department of Health and Human Services: Office of the Inspector General. "Concerns with Manufacturer Rebates"

their bargaining power to split in the surplus gains by demanding rebates that apply when they bear cost in exchange for suppliers setting high drug transaction prices that exploit government risk sharing. Such a split of surplus could explain why bids, which reflect insurer's share of cost, are low while reinsurance is high. These rebate schemes might explain a peculiar pattern of high drug prices at "preferred" and low prices at "non-preferred" pharmacies. More stringent requirements to report rebates at the transaction level and enrollees rebates check, including those for lics payments, may be necessary to temper the strategic.

The 2011 revision to the risk adjustment formula may also be a contributing factor to the recent rise in reinsurance. The new formula indexes risk adjustments to drug prices within the Part D. Insurers gain extra risk adjustment payments for the chronic conditions that experienced drug price increases in prior years. The prospects of extra compensation in the future blunts insurers' incentives to control cost and heightens drug manufacturers' incentives to raise prices. CMS has always been concerned that within-market risk adjustment indexing would diminish incentives to steer patients towards cost-effective treatments Hsu et al. (2010). The point I emphasize is not drug utilization distortions, but rather the effect on drug prices. Out-of-market indexing, as used in 2006-2010 formula, eliminates incentives to manipulate drug prices.

The separate risk adjustments for LIS beneficiaries heighten those incentives to raise prices. Insurers and drug suppliers want to raise drug prices to price discriminate against LIS beneficiaries. Because the risk adjustments separate LIS beneficiaries from regular beneficiaries, the price discrimination inflates risk adjustment factors for LIS beneficiaries without altering risk adjustments for regular enrollees. As risk adjustment factors for LIS beneficiaries ratchet up, the profitability of price discrimination increases, further driving up drug prices. Year over year this pattern would lead to a nearly unbounded spiral of drug prices. The only factor limiting the spiral is the 15% share of cost that insurers bear for catastrophic claims. Drug prices can be kept in check with risk adjustments that group LIS and regular beneficiaries together because insurers bear the full insurance cost for regular beneficiaries and few of them reach catastrophic spending levels.

I suspect 2014 represents an inflection point. There is a lag in updating the risk adjustment model. Between 2011 and 2013, risk adjustments were indexed to drug prices from the years before the revision. Any increases in drug prices in this period would have been made in anticipation of higher risk adjustments in future years. 2014 indexes risk adjustment factors to drug prices in 2011, which is the first year in which the revisions would have had an effect on drug prices. 2014 reinsurance reconciliations could be very high.

The final consideration is the new minimum loss requirement (MLR). Starting in 2014, Part D insurers face a penalty if the ratio of claims paid to revenue falls below 85%. MLRs are intended to limit profits and ensure low premiums. However, as a general statement, a binding MLR limits incentives to keep costs—claims paid—low. Reinsurance may be accentuating that incentive problem. The guidelines for MLR count bids and prospective payments for reinsurance as part of revenue $b_j + reins_j$. Table 7 reports average revenue in Part D as $70 + 60 = 130$ for 2015. Costs include the insurer share of claims cost in the initial coverage zone (75%) and catastrophic (15%) and reinsurance (80%). To comply with MLR for this example, these three components cannot exceed $.85 * 130 = 111$, which limits the profit markup to $19. The insurer does not actually bear a cost anywhere near $111 because the government pays reinsurance claims costs. If the reinsurance component of the bid is removed from the cost and revenue component of the calculation, an 85% MLR applied to the $70 bid limits the profit markup to $10. Notice the profit margin is much higher when reinsurance factors into the equation. To increase allowable profits margins under MLRs insurers want to drive up reinsurance while keep their own cost in check. They do so by price discrimination against reinsurance claims.

There are winners and losers from these trends. There is little effect on regular beneficiaries because premiums remain stable and insurers can shield them from high point-of-sale drug prices by setting low copays. Market prices for drugs and premiums don't matter for LIS beneficiaries because they are so heavily subsidized. Drug suppliers gain extra surplus by profiting off of reinsurance. There are ways for insurers to share in this surplus too. The clear losers are the tax payers financing the rising government payments. There may be collateral damage in other health care markets if rising drug prices induced by Part D increase costs for employer sponsored health insurance, Medicaid, private payers, and ACA exchange health plans. Policy makers should consider eliminating or reducing reinsurance so that insurers have an incentive to lower cost, not raise cost. Program subsidies could then be directed towards the least distortionary subsidy of all, the general subsidy (λ_t). It has already fallen from 0.65 to 0.52 but would return to 0.745 if reinsurance is eliminated.

7 Conclusion and Policy Recommendations

In this paper, I examine distortions in the Medicare Part D market stemming from the low income subsidy (LIS) program rules and its link with the three "Rs" in the risk adjustment mechanism. The LIS program provides premium and cost sharing subsidies for low income beneficiaries and Medicare/Medicaid dual eligibles, over and above those available to regular enrollees. LIS beneficiaries compose about 20% of the Medicare population and have substantially higher drug expenditures. At upper tail, 20% of LIS enrollees reach catastrophic spending levels, only 3% of regular enrollees reach those level. There are two key features of the LIS program that act to distort the market. Premium and copay subsidies and the automatic and random assignment of dual eligibles who do not actively select a plan. Subsidies and auto-enrollment distort residual demand elasticities and create a demand discontinuity at the market price indexed threshold at which plans become eligible to receive auto-enrollees. This induces a bunching of prices at the threshold and sequesters LIS enrollees into a limited set of plans. I estimate a structural demand/supply model to illustrate that the choice friction of auto-assignment generates a large demand discontinuity. The threshold serves its purpose of acting as price ceiling to prevent prices from rising in a market with otherwise highly inelastic demand due to subsidies. Profit markups for LIS eligible plans are slightly above that for higher cost non-LIS plans and lower than markups on enhanced plans. Despite the apparent success of the LIS threshold to temper prices, the indexing scheme to past market shares inducing a dynamic cycling pattern whereby insurers oscillate prices above and below the threshold. This may have adverse effects on patients who are reassigned to another plan and face the costs associated with coverage disruptions. A weighting scheme that excludes LIS beneficiaries and a random assignment scheme that distributes enrollees in proportion to price and coverage benefits could eliminate cycling and improve incentives to compete for LIS auto-enrollees. Alternatively, financial incentives could be used induce LIS beneficiaries to opt of random assignment and make their own plan decision. Financial incentives could be quite effective because the demand estimates show the low income segment of the population has very elastic demand.

I next consider the three elements of the risk adjustment mechanism (risk adjustments, risk corridors, reinsurance). Payments for all three components have exhibited great imbalance and have been highly volatile throughout the program's history. The volatility has prompted CMS to indefinitely continue risk corridors despite the intent to phase them out in 2008. A troubling trend in reinsurance has taken shape in the past few years. Reinsurance, government risk sharing of claims costs for beneficiaries reaching catastrophic spending levels, has more than tripled since 2007 and accounts for about half of the cost of insurance. Moreover, insurers prospective predictions of reinsurance spending chronically underestimate the actual cost. That gap is growing. Paralleling these trends, premiums remain stable and total payments to insurers for their share of cost have been falling and persistently beat expectations.

I begin the analysis of risk sharing by providing theory and evidence from risk sharing payment data to show how the LIS program contributes to the volatility. Risk adjustments are transfer payments between plans that are intended to prevent adverse selection problems. In principle, insurers should be indifferent as whether they have a low or high cost pool of enrollees. Risk corridors are a risk sharing scheme with the government to provide insurance to insurers in case their costs do not align with risk adjustments. Despite the best intentions for risk adjustments to accurately predict costs, insurers are able to cream skim beneficiaries with drug expenditures lower than that predicted by the risk adjustment formula through sophisticated formulary benefit design practices. However, cream-skimming doesn't work on LIS auto-enrollees whose blunted choice incentives renders them unresponsive to cream-skimming efforts. This results in an evenly distributed pool of risk across LIS eligible plans. In contrast, cream-skimmed regular enrollees are less costly on a risk adjusted basis than non cream-skimmed auto-enrollees. The segmentation of LIS enrollees into LIS eligible plans explains the variation in risk corridor payments across insurers, and the within-insurer across year volatility that arises due to the yearly cycling of plans in and out of LIS eligibility. The discord in selection further distort pricing by creating a cost discontinuity at the LIS threshold. The combination of a demand and cost discontinuity induces insurers to bunch their bids right above the threshold, in effect turning the threshold into a price floor. Insurers have found a way to circumvent the price floor to attract more regular enrollees by offering a slightly enhanced plan priced lower than basic plans. These plans are not substantially different from basic plans, but their designation as "enhanced" allows them to avoid unfavorably selected LIS enrollees. There are proposals to eliminate these low priced enhanced plans. However, their elimination would only take away meaningful competition. Other efforts to properly score risk would be more effective and preserve competition. The 2011 revision to the risk adjustment formula, that indexes to prior year Part D claims and separates LIS and regular beneficiaries was intended to correct these discrepancies. I provide evidence that the revision was effective to even risk adjusting of LIS and regular enrollees, but it may have unleashed more severe problems related to reinsurance.

I next consider reinsurance. There is similar pattern in reinsurance reconciliation payments (the difference between actual and predicted reinsurance) that indicate LIS beneficiaries are unfavorably selected relative to risk adjustment predictions. However, the 2011 revision did not eliminate the discrepancies as was the case for risk corridors. This suggest something else is causing an excessive level of reinsurance being paid to LIS eligible plans.

The fundamental problem with government risk sharing is that insurers have no incentive to control cost when the government bears cost. I examine over 2,000,000 pharmacy transactions and show that insurers and drug suppliers are able to employ sophisticated drug price discrimination schemes to raise prices on drug claims in which the government bears cost (through reinsurance and low income cost sharing) relative to prices when the insurer or enrollee bears costs. They can price discriminate against the drugs more heavily used by LIS beneficiaries, and even for the same drugs depending on formulation, strengths, days-supply, and pharmacy outlet. Insurers can price discriminate across plans, setting high prices for LIS eligible plans and low prices for non-LIS eligible plans. They exploit seasonality to price higher later in the year when more beneficiaries have entered catastrophic spending levels. The estimates may understate the magnitude of price discrimination. Insurers, manufacturers, drug suppliers, and pharmacies can use rebates, not-applied to transactions, to maximize the extraction of rents from reinsurance and split surplus. It's conceivable that price discrimination has grown so strong to have triggered the astronomical price increases plaguing generic drug markets. These patterns Anti-trust authorities should closely consider drug price setting when reviewing horizontal mergers case amongst insurers and

Recent changes to the program may have heightened these price discrimination incentives and could be leading to an unraveling of the market. The major changes to the risk adjustment model (indexing to in-market pricing, segregation of LIS enrollees) promote higher drug prices because risk adjustment payments rise in future years to compensate for higher prices. Insurers and drug suppliers would target those rises towards LIS beneficiaries because the government bears a large share of their cost. Despite some of the imbalances in the 2006-2010 risk adjustment formula, it provides better incentives to temper drug price discrimination by integrating LIS and regular beneficiaries and indexing to outside market drug prices. The new 2014 minimum loss ratio requirements may further promote price discrimination against reinsurance because the accounting practices factor those payments into revenue, even though they are not costs to insurers, rather pass-through payments.

Risk sharing and LIS program act to distort market outcomes in quite complicated and interconnected ways. Patchwork solutions to specific problems, such as re-index the LIS threshold, recalibrating risk adjustments, capping premiums may not be the best solution to solve the overarching issues, and could actually be amplifying problem. More drastic reform may be necessary. Congress has proposed legislation mandating prices controls for LIS and reinsurance claims so that the government can take back control of the drug prices it finances. Alternatively, policy makers should consider eliminating or at least reducing its role in risk sharing for reinsurance and risk corridors. Much of the ability for insurers to price discriminates can be attributed to the plan assignment rules that segregate LIS and regular enrollees. The fix is either more integration or a complete separation. LIS enrollees could be integrated with regular enrollees by uniformly distributing then across all plans or in proportion to regular beneficiary enrollment. The benefits of reducing drug price discrimination might outweigh any adverse effects on premium competition. At the other extreme, policy makers should consider removing LIS beneficiaries, in particular dual eligibles, from Part D markets and placing them in their own program, possibly a return to Medicaid.

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.

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

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. *American Economic Review*, 104(10):3335–3364, 2014.

C. Carey. Government Payments and Insurer Benefit Design in Medicare Part D. *Michigan working paper*, 2014.

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

F. Decarolis, M. Polyakova, and S. Ryan. Welfare Effects of Supply-Side Regulation in Medicare Part D. *Stanford working paper*, 2015.

M. Duggan and F. Scott-Morton. The Effect of Medicare Part D on Pharmaceutical Prices and Utilization. *The American Economic Review*, 100(1):590–607, 2010.

K. Ericson. Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange. *American Economic Journal: Economic Policy*, 6(1):38–64, 2014.

J. Glazer and T. McGuire. Optimal Risk Adjustment in Markets with Adverse Selection: an Application to Managed Care. *The American Economic Review*, 90(4):1055–1071, 2000.

B Handel. Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *The American Economic Review*, 103(7):2643–2682, 2013.

J. Hartigan and P. Hartigan. The Dip Test of Unimodality. *The Annals of Statistics*, pages 70–84, 1985.

F. Heiss, A. Leive, D. McFadden, and J. Winter. Plan Selection in Medicare Part D: Evidence from Administrative Data. *Journal of Health Economics*, 32(6):1325–1344, 2013.

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. Hsu, V. Fung, J. Huang, M. Price, R. Brand, R. Hui, B. Fireman, W. Dow, J. Bertko, and J. Newhouse. Fixing Flaws in Medicare Drug Coverage that Prompt Insurers to Avoid Low-income Patients. *Health Affairs*, 29(12):2335–2343, 2010.

J. Ketcham, C. Lucarelli, E. Miravete, and M. Roebuck. Sinking, Swimming, or Learning to Swim in Medicare Part D. *The American Economic Review*, 102(6):2639–2673, 2012.

P Klemperer. Competition When Consumers Have Switching Costs: An Overview With Applications to Industrial Organization, Macroeconomics, and International Trade. *The Review of Economic Studies*, 62(4):515–539, 1995.

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

B. Madrian and D. Shea. The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187, 2001.

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

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*, 2012.

D. Miller and J. Yeo. The Consequences of a Public Health Insurance Option: Evidence from Medicare Part D. *Working Paper Clemson University, Singapore Management University*, 2013.

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

M. Polyakova. Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D. *Stanford working paper*, 2015.

B. Silverman. Using Kernel Density Estimates to Investigate Multimodality. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 97–99, 1981.

L. Summer, J. Hoadley, and E. Hargrave. The Medicare Part D Low-income Subsidy Program: Experience to Date and Policy Issues for Consideration. *Kaiser Family Foundation*, 2010.

L. Summer, J. Hoadley, and E. Hargrave. Medicare Part D in its Ninth Year: The 2014 Marketplace and Key Trends, 2006-2014. *Kaiser Family Foundation*, 2014.

E Tamer. Incomplete Simultaneous Discrete Response Model with Multiple Equilibria. *The Review of Economic Studies*, 70(1):147–165, 2003.

Y. Zhang, C. Zhou, and S. Baik. A Simple Change to the Medicare Part D Low-Income Subsidy Program could save $5 Billion. *Health Affairs*, 33(6):940–945, 2014.

8 Appendix: Subsidy Elasticities

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 market size term M_{jmt} depends on whether the plan's basic premium is above or below the LIS threshold $\bar{s}^{LIS}_{mt}$. By substituting in the subsidy rules given in equations 15 and 19 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'.[18] 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, which requires calculating 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

$$\begin{aligned}
\eta_{jjmt} &= \frac{\partial s_{jmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{jmt}} = -\alpha_i b_{jmt}\left[(1-s_{jmt}) - \kappa_i\left(1-\frac{\gamma_{jmt}}{1+\gamma_{jmt}}\right)(1-s_{jmt}) + \kappa_i \frac{w^{lis}_{jmt}}{1+\gamma_{jmt}} s^{above}_{mt} - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjmt} &= \frac{\partial s_{kmt}}{\partial b_{jmt}}\frac{b_{jmt}}{s_{kmt}} = -\alpha_i b_{jmt}\left[-s_{jmt} + \kappa_i\left(1-\frac{\gamma_{jmt}}{1+\gamma_{jmt}}\right)s_{jmt} + \kappa_i \frac{w^{lis}_{jmt}}{1+\gamma_{jmt}} s^{above}_{mt} - (1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jmt}}\tilde{w}_{jmt-1}s_{0mt}\right] \\
\eta_{kjm't} &= \frac{\partial s_{kmt}}{\partial b_{jm't}}\frac{b_{jm't}}{s_{kmt}} = -\alpha_i b_{jm't}\left[-(1-\kappa_i)\frac{\lambda_t}{1+\gamma_{jm't}}\tilde{w}_{jm't-1}s_{0mt}\right]
\end{aligned} \quad (31)$$

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

[18]Because the weights $\tilde{w}_{jmt-1}$ are based on lagged enrollment, I could also calculate cross price elasticities across time. I do not because the 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, $\bar{s}_{mt}^{LIS}$. The intuition is that when a plan increases its bid, it raises the maximum subsidy amount $\bar{s}_{mt}^{LIS}$. 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 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

$$\begin{array}{rclcl}
\eta_{jjmt} & = & \frac{\partial s_{jmt}}{\partial b_{jmt}} \frac{b_{jmt}}{s_{jmt}} & = & -\alpha_i b_{jmt} \left[(1 - s_{jmt}) - \kappa_i \frac{w_{jmt}^{lis}}{1+\gamma_{jmt}} (1 - s_{mt}^{above}) - (1 - \kappa_i) \frac{\lambda_t}{1+\gamma_{jmt}} \tilde{w}_{jmt-1} s_{0mt} \right] \\
\eta_{kjmt} & = & \frac{\partial s_{kmt}}{\partial b_{jmt}} \frac{b_{jmt}}{s_{kmt}} & = & -\alpha_i b_{jmt} \left[-s_{jmt} - \kappa_i \frac{w_{jmt}^{lis}}{1+\gamma_{jmt}} (1 - s_{mt}^{above}) - (1 - \kappa_i) \frac{\lambda_t}{1+\gamma_{jmt}} \tilde{w}_{jmt-1} s_{0mt} \right] \\
\eta_{kjm't} & = & \frac{\partial s_{kmt}}{\partial b_{jm't}} \frac{b_{jm't}}{s_{kmt}} & = & -\alpha_i b_{jm't} \left[-(1 - \kappa_i) \frac{\lambda_t}{1+\gamma_{jm't}} \tilde{w}_{jm't-1} s_{0mt} \right]
\end{array} \tag{32}$$

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.

Furthermore, I must account for automatic enrollment which is determined by the bid. Recall, a plan qualifies for automatic enrolles if $p_{jmt}^{basic} \leq \bar{s}_{mt}^{LIS}$ 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.

9 Appendix: Estimating Marginal Cost

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 is not 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 I 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 an inversion of 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

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

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.

$$mc_{jmt}^{basic} = mc_{kmt}^{basic} \tag{33}$$

Consider an example for a plan that offers an enhanced plan (plan 1) and a basic plan (plan 2). The first order condition of equation 10 with respect to the bids is:

$$\begin{aligned} 0 &= s_1 + (b_1 - mc_1)\frac{\partial s_1}{\partial b_1} + (b_2 - mc_2)\frac{\partial s_2}{\partial b_1} && (34) \\ 0 &\geq s_2 + (b_1 - mc_1)\frac{\partial s_2}{\partial b_1} + (b_2 - mc_2)\frac{\partial s_2}{\partial b_2} && (35) \end{aligned}$$

If the basic plan 2 is priced away from the LIS 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.[19] 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

$$0 = s_1 + (b_1 - mc_1)\frac{\partial s_1}{\partial b_1} + \left(b_2 - \frac{mc_1}{1+\gamma_1}\right)\frac{\partial s_2}{\partial b_1} \tag{36}$$

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

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

$$\mathbf{0} = \mathbf{s} + \Delta(\mathbf{b} - \mathbf{mc}) \tag{37}$$

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

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

cross-price elasticities. Marginal cost can be solved for by inverting the system of equations:

$$\mathbf{mc} = \mathbf{b} + \Delta^{-1}\mathbf{s} \tag{38}$$

For firms with plans priced at the threshold the first order conditions are modified by imposing the cost restriction in (33):

$$R\mathbf{0} = R\mathbf{s} + R\Delta\mathbf{b} - R\Delta R'R\mathbf{mc} - R\Delta' M_\gamma R'R\mathbf{mc} \tag{39}$$

The restriction matrix R has dimension $(N - N_{\bar{s}^{LIS}} \times N)$ where $N_{\bar{s}^{LIS}}$ is the number of plans priced at the threshold. The jj entry is a one for the first $j = 1, \ldots N - N_{\bar{s}^{LIS}}$ 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 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 $R\mathbf{s}$.

$$R\mathbf{mc} = (R\Delta R' + R\Delta' M_\gamma R')^{-1}(R\mathbf{s} + R\Delta\mathbf{b}) \tag{40}$$

The remaining $N_{\bar{s}^{LIS}}$ 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. 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 offer one in just a few. 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. I choose the enhanced plan with observed product characteristics closest in characteristic space to the basic plan. The third 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 price within a few dollars of the threshold. Including incomplete information greatly complicates the model. Instead, I designate any plan that prices within a small dollar range ($2) of the threshold as being a threshold plan. The fourth, most important limitation, is multiple equilibrium (Tamer, 2003). This framework cannot predict which plans do and do not enter below the threshold. This issue prevents estimating a full equilibrium supply-side model. Instead I estimate the model for 1 firm, taking as fixed the bids of all other insurers.

10 Appendix: Bids Around LIS Threshold

The theory predicts bunching of bids both above and below the LIS threshold. Because of imperfect information about the location of the threshold, the mode of bids is bounded away from the threshold, which should result in 2 distinct modes centered at the threshold. Figure 10 shows kernel density plots of bids in a local neighborhood of the threshold for plans that did not have LIS status in the prior year. Figure 11 depicts for plans that had LIS status in the prior. Multimodality, and in particular the dip at the LIS threshold, can be observed for both types of plans in most years. It is most apparent for prior year non-LIS plans and for the years 2007 through 2009. In 2011 and 2012, the densities appear more unimodal. These years follow the implementation of CMS' updated risk scoring model that factors in LIS status. The interpretation of these figures is complicated by the deminimis rules. They were effective 2007, 2008, discontinuity 2009,2010, and reintroduced thereafter, albeit with different rules. In the effective years, there appears to be a third mode (second dip) above the LIS threshold, that is not present in 2009, 2010.

One should be careful drawing inference about modality from kernel density plots because of sensitivity to smoothing parameters. As an example, the 2011 prior year LIS figure has many modes as drawn. The extra modes could be a spurious result of smoothing. However they could also reflect non-spurious differences in the distributions of bidders' uncertainty about the LIS threshold location. The latter explanation does not discredit the hypothesis of bunching above and below. Unfortunately, there is little power to distinguish in such small samples. Table 13 reports Hartigan and Hartigan (1985) dip test statistics and p-values. The null is unimodality. The largest dips are found for prior year non-LIS plans, in particular 2009. The 2011 and 2012 dip strongly rejects bimodality.

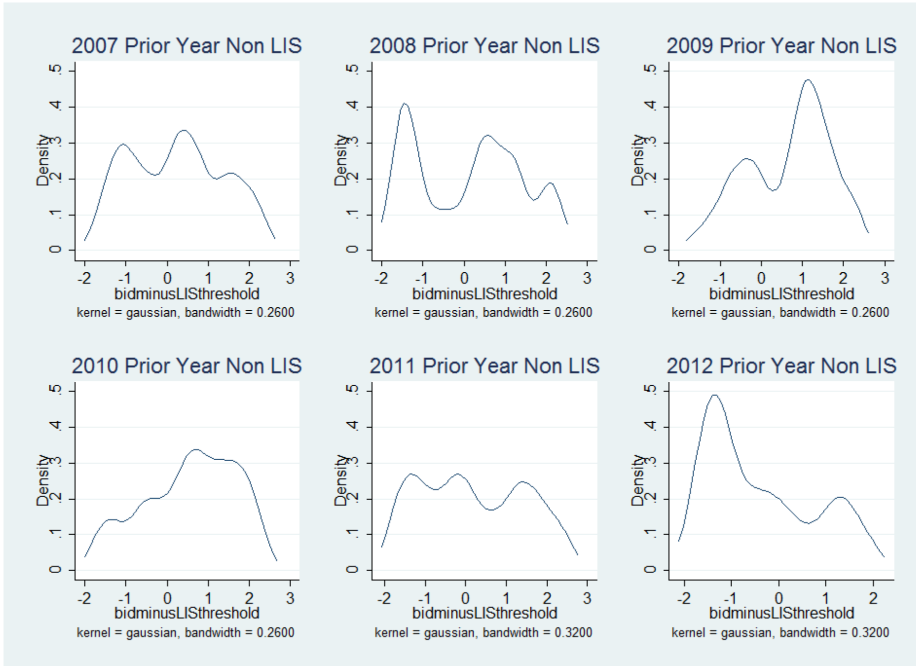

Figure 10: Bids: Prior year non-LIS plans

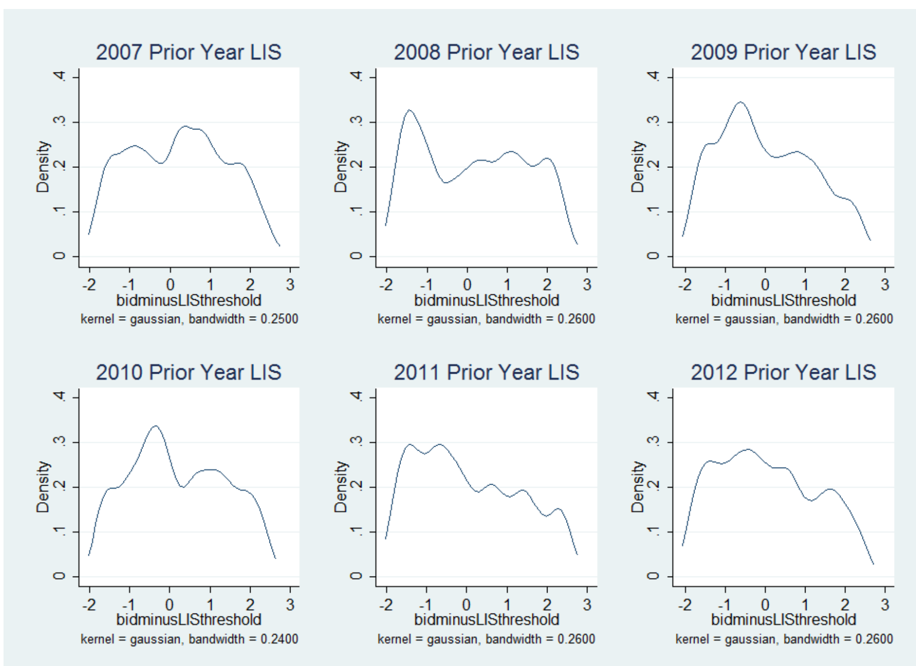

Figure 11: Bids: Prior year LIS plans

Table 13: Dip Test for Multimodality of Bids Around LIS Threshold

	Prior year non-LIS			Prior year LIS		
	obs	dip	p-value	obs	dip	p-value
2007	86	0.045	0.346	124	0.041	0.269
2008	35	0.086	0.143	218	0.031	0.270
2009	37	0.101	0.052	159	0.033	0.357
2010	113	0.051	0.131	144	0.041	0.208
2011	148	0.039	0.227	124	0.031	0.645
2012	61	0.045	0.576	136	0.026	0.848

The range is restricted from $1.80 below the threshold to 2.50 above the threshold.