Regulation of Public Health Insurance

by

Maria A. Polyakova

B.A. Economics and Mathematics, Yale University (2008)
B.A. German Studies, Yale University (2008)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

Signature redacted

Author .....
Department of Economics
May 15, 2014

Signature redacted

Certified by .....
Amy Finkelstein
Ford Professor of Economics
Thesis Supervisor

Signature redacted

Certified by .....
Nancy Rose
Charles P. Kindleberger Professor of Applied Economics
Thesis Supervisor

Signature redacted

Accepted by .....
Michael Greenstone
3M Professor of Environmental Economics
Chairman, Departmental Committee on Graduate Studies

Regulation of Public Health Insurance

by

Maria A. Polyakova

Submitted to the Department of Economics
on May 15, 2014, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The first chapter takes advantage of the evolution of the regulatory and pricing environment in the first years of a large federal prescription drug insurance program for seniors - Medicare Part D - to explore interactions among adverse selection, switching costs, and regulation. I document evidence of both adverse selection of beneficiaries across contracts and switching costs for beneficiaries in changing contracts within Medicare Part D. Using an empirical model of contract choice and contract pricing, I show that in the present environment, on net, switching costs help sustain an adversely-selected equilibrium with large differences in risks between more and less generous contracts. I then simulate how switching costs may alter the impact of "filling" the Part D donut hole as implemented under the Affordable Care Act. I find that absent any switching costs, this regulation would have eliminated the differences in risks across contracts; however, in the presence of the switching costs that I estimate, the effect of the policy is largely muted.

The second chapter (co-authored with Francesco Decarolis and Stephen Ryan) explores federal subsidy policies in Medicare Part D. We estimate an econometric model of supply and demand that incorporates the regulatory pricing distortions in the insurers' objective functions. Using the model, we conduct counterfactual analyses of what the premiums and allocations would be in this market under different ways of providing the subsidies to consumers. We show that some of the supply-side regulatory mechanisms, such as the tying of premiums and subsidies to the realization of average "bids" by insurers in a region, prove to be welfare-decreasing empirically.

The third chapter studies two competing systems that comprise the German health insurance landscape. The two systems differ in the ability of insurers to underwrite individual-specific risk. In contrast to the community rating of the statutory insurance system, enrollees of the private plans face full underwriting and may be rejected by the insurers. I empirically assess to what extent the selection of "good risks" dominates the interaction between the two systems, using a regression discontinuity design based on statutory insurance enrollment mandates. I do not find compelling evidence of cream-skimming by private insurers from the statutory system. Motivated by this finding, I quantify the change in consumer welfare that would result if the government relaxed the statutory insurance mandate to lower income levels.

Thesis Supervisor: Amy Finkelstein
Title: Ford Professor of Economics

Thesis Supervisor: Nancy Rose
Title: Charles P. Kindleberger Professor of Applied Economics

Contents

[1] **This Chapter is co-authored with Francesco Decarolis and Stephen Ryan.** Decarolis is grateful to the Sloan Foundation (grant 2011-5-23 ECON). We also gratefully acknowledge support from the NSF (SES-1357705)

Acknowledgments

First and foremost, I am deeply grateful to Amy Finkelstein and Nancy Rose for having been overwhelmingly dedicated thesis advisors. Without their enthusiasm, wisdom, humor and patience, this thesis would not have been possible. I cannot thank them enough for the endless hours they spent talking to me about my research, reading my drafts, writing detailed comments, and encouraging me to dig deeper and not give up when initial results were not perfect. Amy has been the inspiration of my interest in social insurance at large, and health insurance in particular, while Nancy introduced me to the fascinating world of regulation. In their different ways, they taught me a lot about life in and outside of an academic career, and I can only aspire to their caring, compassionate, and inspiring examples in my own research and teaching. I also thank Stephen Ryan, who has introduced me to the world of structural econometric models, which, as he put it with his unbounded enthusiasm, could be used to analyze "awesome big picture questions!" I am grateful for his mentorship and positive energy throughout my years as his student and co-author.

I am also grateful to a long list of other faculty at MIT, who were generously available to talk about my projects and gave invaluable advice. Jonathan Gruber often brought me back to think about the big picture and policy implications, while Jim Poterba gave very helpful suggestions of how to talk about my research on the job market, all while giving exactly the right words of encouragement during that stressful time. I thank Glenn Ellison, Sara Fisher Ellison, Anna Mikusheva, Panle Jia, David Autor, Bob Gibbons, Michael Greenstone, Paulo Somaini, and Heidi Williams, who all contributed to this thesis through their helpful comments and suggestions.

If I had to name one person, who was responsible for sparking my original interest in academic economics, that person would certainly be my undergraduate advisor - Tim Guinnane. I took Tim's economic history class as a sophomore in college. Upon completing that class, I discovered two things - first, that I wanted to teach economics just like Tim, and second, that I did not know how to properly use articles in English. As I complete this thesis and start my academic career, I can claim some progress on the first front; but, unfortunately, hardly any on the second. Tim's continued mentorship throughout my time in graduate school has been invaluable and I will always be grateful.

This thesis would not have been possible without the friendly, creative and fun atmosphere that my classmates at MIT are responsible for. I thank Dan Rees and Brad Shapiro for being wonderful first-year roommates, with whom I could share dinner, homework and exam anxiety. I thank Patricia Gomez-Gonzalez, Annalisa Scognamiglio, Stefanie Stantcheva, Dana Foarta, and Nina Harari for the "women's dinners" and for always being there for deep conversations and moments of self-doubt. Adam Sacarny, Ben Feigenberg, Conrad Miller, Isaiah Andrews, Miikka Rokkanen, Anil Jain, Juan Passadore, Anton Tsoy, and Mike Yankovich, have all made these years special. Outside of MIT, I am grateful to Svetlana Dotsenko, Yelena Kosheleva, and Jing Xia, who made sure my sanity was kept in check.

Last, but far from least, I am indebted to my family for their love and support throughout my time in graduate school. I am grateful to my parents and grandparents, who are responsible for cultivating my affinity to analytic inquiry. I thank my father for only mildly criticizing the fact that I planned on becoming an academic, for considering the topic of my dissertation only mildly useless, and then spending hours enthusiastically debating with me on how a country should set up its healthcare system. I thank my mother for making sure the debate above did not continue till dawn. Finally, I am infinitely grateful to my husband without whose unconditional love, patience, innumerable long-haul flights to Boston, cooking, talking-about-insurance-24/7, and LaTex knowledge, this thesis would not have been possible, both literally and figuratively.

List of Tables

List of Figures

Chapter 1

Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D

1.1 Introduction

Outsourcing a public health insurance benefit to private insurers creates a familiar trade-off between potential efficiency gains from competition and potential efficiency losses from adverse selection. We know much less, however, about the role that other market imperfections, such as switching costs, play in competitive insurance settings. In this paper, I utilize the institutional environment of Medicare Part D prescription drug insurance program to empirically analyze how switching costs affect the re-distribution of risks across contracts in response to market evolution, as well as how they alter the consequences of regulatory interventions that directly change the contract space.

Medicare Part D is a large insurance program introduced in 2006 that currently enrolls approximately 32 million beneficiaries, with annual federal spending of around $63 billion. This heavily subsidized benefit is administered entirely by private insurers that are extensively regulated. The insurers offer a variety of plans, making individual contract choice a prominent feature of the program.

Motivated by the evolution of the regulatory and pricing environment of Medicare Part D, I use the detailed administrative records on individual-level enrollment choices, ex ante health

[0]I am greatly indebted to Amy Finkelstein, Nancy Rose, Stephen Ryan, and Jonathan Gruber for their invaluable guidance and support throughout this project. I thank Jason Abaluck, Isaiah Andrews, Andreas Breiter, Francesco Decarolis, Glenn Ellison, Sara Fisher Ellison, Kate Easterbrook, Michael Greenstone, Nathan Hendren, Gaston Illanes, Ray Kluender, Brad Larsen, Anna Mikusheva, Manisha Padi, Maxim Pinkovskiy, Jim Poterba, Adam Sacarny, Brad Shapiro, Annalisa Scognamiglio, Stefanie Stantcheva, Michael Whinston, Joachim Winter and participants at the MIT Public Finance and Industrial Organization field lunches and workshops for very helpful comments, suggestions and discussions. I also thank Jean Roth at NBER and Sarah Brunsberg at ResDAC for their help in navigating the Medicare data, as well as the Max Planck MEA for their hospitality.

risk, and realized ex post spending in the program to explore how switching costs interact with adverse selection and with the regulatory interventions that change the relative generosity of plans. The empirical analysis shows that the direction in which switching costs affect adverse selection depends crucially on how both the relative prices and the relative generosity of contracts evolve in comparison to the initial conditions. Taking the example of contract changes stemming from Medicare Part D's minimum standard policy known for its gap in coverage, or "donut hole", the paper also demonstrates the implications of the interaction between adverse selection and switching costs for policy instruments designed to ameliorate adverse selection.

The paper is structured as follows. The key pieces of Part D's institutional environment that are necessary for the subsequent empirical analysis are summarized in Section 1.2. Section 1.2 also describes the sample of the nationwide administrative records from the Centers of Medicare and Medicaid used in this paper. In Section 1.3, I turn to the descriptive analysis of the data.

The data suggest that despite numerous regulatory provisions designed to mute adverse selection in this market, both cross-sectional and dynamic selection exists in Medicare Part D. Surprisingly, given the extent of policy attention to this concern, adverse selection in Medicare Part D has received little attention in the empirical literature.[1] To fill this gap, I document several pieces of evidence for the presence of adverse selection. First, I present the results of the cross-sectional positive correlation tests as described in Chiappori and Salanie (2000). These tests suggest a large degree of asymmetric information: the most generous plans attracted individuals whose annual drug spending was more than a standard deviation above the spending in the least generous plans. Although it is typically hard to disentangle selection from moral hazard in the positive correlation tests, Medicare Part D setting provides me with a way of addressing this problem. I use a rich set of moral hazard-free ex ante diagnostic information about the beneficiaries to recompute the positive correlation tests, isolating the lower bound of the cross-sectional adverse selection. These tests reveal that the most generous contracts in Medicare Part D attracted individuals with substantially higher (about $730 a year) expected spending. Consistent with these cross-sectional results is also the evidence of dynamic adverse selection, which we see in the rapid unraveling of the contracts that offered full coverage in the "donut hole."

I then turn to the question of whether Medicare Part D exhibits switching costs, which may be an important factor in understanding the development over time of the risk-allocation patterns highlighted above. While the literature has focused on the "inconsistency" of individual choices relative to a perfectly rational benchmark (e.g. Abaluck and Gruber 2011, 2013; Heiss et al. 2013; Ketcham et al. 2012), I do not assess the efficiency of choices here; rather, I document that, whichever choices individuals make, changing their initial choices appears costly.

Identifying switching costs or "true" state dependence separately from unobserved heterogeneity is a notoriously difficult problem (e.g., Heckman, 1991). Medicare Part D's institutional setting and administrative records, however, provide nearly ideal variation and data for documenting the

[1]To the best of my knowledge, no systematic analysis of selection in the spirit of the established literature testing for asymmetric information has been conducted in this important setting.

presence of structural state dependence. First, there is substantial variation in the choice set that is available in the program over time. Second, every year, new beneficiaries enter the program, as they become eligible for Medicare upon turning 65. Using this variation, I show that choices of the "continuing" cohorts persistently reflect the market conditions of the year in which these individuals made their first choices, while the choices of the otherwise similar newly entering cohorts are different and reflect the current market conditions. This approach to documenting switching costs is in the spirit of Handel (2013)'s approach to documenting inertia in employer-provided health insurance. Consistent with the cross-cohort choice comparison, I also find that existing cohorts are less price-sensitive than the newly entering cohorts. These findings provide micro-level identification for the presence of switching costs in Part D and thus support the hypothesis made in the earlier literature based primarily on aggregate data (Ericson, 2013) that inertia is an important feature of this market.[2]

With the reduced-form evidence of adverse selection and switching costs in hand, I proceed to a utility-based choice model of the individual preferences for insurance plans in Section 1.4. I estimate a flexible choice model with heterogeneous preferences and private information about health risk that allows me to quantify switching costs, adverse selection, and the individuals' willingness to pay for different contract features conditional on their risk. The estimation results, discussed in Section 1.5.1, suggest that switching costs are large and critical for explaining enrollment paths over time. The estimates further suggest that information about risk plays an important role in determining individual choices, which is consistent with the descriptive evidence that the self-selection of beneficiaries in the Part D environment leads to adverse selection. The model fits the data well and correctly captures the magnitude of the key descriptive patterns.

Counterfactual analyses in Sections 1.5.2 and 1.5.3 utilize the estimated choice model to explore the complex interactions among adverse selection, switching costs, and minimum standard policy regulations in the Medicare Part D setting. The textbook model of adverse selection in insurance hinges on the idea that costlier individuals will choose the most comprehensive plan; this plan will experience high costs and will have to increase its premium. In turn, the least costly individuals will drop out of the plan and the spiral will continue unraveling. It then seems intuitive that if we introduce a switching friction into this model, it may stop the unraveling process right at the first loop: if the friction is sufficiently high, no individuals will leave the most comprehensive plan despite the increase in its price. Indeed, Handel (2013) provides an empirical example of switching costs restricting the selection spiral in an employer-provided health insurance setting. In that setting, the relative price of a more generous contract increases over time, but individuals do not re-adjust their choices due to switching costs. This evidence is sometimes interpreted as confirming the intuition that switching costs *must* be helpful in ameliorating adverse selection.[3] In general, however, the

[2]Since completing this paper, I learned of contemporaneous work in Ho, Hogan, and Scott Morton (2013), which documents that switching is rare in Part D and analyzes channels that may be driving switching, as well as the implications of these market features for supply-side behavior.

[3]Handel (2013), however, points out that the direction of the results is specific to the considered environment and enrollee population.

direction of the effect of switching costs on selection is ambiguous. If the relative price of the more comprehensive contract fell instead of increasing over time, the result would have likely been the opposite.

In practice, the latter possibility of relative premiums falling is especially relevant in more complex insurance settings, such as public health insurance exchanges. In such settings, a large number of competing plans are differentiated both vertically and horizontally, their prices and characteristics are affected by a host of regulatory interventions, and, due to complex subsidy provisions and strategic pricing, the relative premiums of plans may not perfectly reflect the actuarial differentials among them. The combination of these factors implies that in such markets the relative price of a more comprehensive contract, for example, may go down rather than up, or its characteristics may adjust relative to the other contracts in a way that favors less acute sorting. In such cases, the switching cost would have the unintuitive effect of exacerbating the adverse selection concern. In Section 1.5.2, I provide two stylized examples using a significantly simplified, yet realistic, version of the Part D contract space that illustrate the opposite effects that switching frictions may have on selection. In the same section, I proceed to describe the full counterfactual simulation that considers all contracts offered in Medicare Part D. In this simulation I find that Part D market evolved in a way that made switching costs support, rather than mute, adverse selection. Specifically, I find that without the switching friction, the difference in the average risk between the least and the most generous contracts would have been 21% lower.[4]

A common channel driving the changes in the relative prices and generosity of contracts in public health insurance settings is regulatory intervention. The fact that switching costs alter the response of risk-sorting to contract changes implies that switching costs will also significantly alter how policy instruments used by the government to regulate this market will impact the allocation of risk. In Section 1.5.3, I explore this hypothesis on the example of Part D's minimum standard regulation. This regulation - the so-called Standard Defined Benefit (SDB) - specifies the minimal amount of coverage that the participating insurers have to offer in their plans. Considering this particular regulation is highly policy-relevant, as the provisions of the Affordable Care Act (ACA) envision significant expansion of the minimum benefit to eventually eliminate the "donut hole". My simulations suggest that switching costs significantly mute the ability of the minimum standard regulation to change the distribution of risks across contracts. For example, I estimate that without switching costs, the expansion of the minimum standard by "filling" the donut hole would substantially reduce adverse selection; with switching costs, however, this expansion has little effect on sorting. In the spirit of the theory of second-best, these results demonstrate the importance of accounting for the interaction among different market failures in health insurance markets and how the correction of one market failure, such as demand-side frictions, may change the effect that regulatory instruments targeted at correcting a different market failure, such as

[4]These simulations allow the insurers to respond to the adjustments in the risk allocation induced by costless switching by adjusting their premiums, but not other contract characteristics. The premium adjustments follow the contract pricing model discussed in Section 1.4.2. Since price adjustments follow lagged risk-based pricing, they tend to amplify the effects that removing the switching costs has on the allocation of risks.

adverse selection, have on the market.

The analysis in this paper is related to several literatures. First, the paper is related to the growing body of literature that analyzes the Medicare Part D program. Most of this literature has focused on assessing the rationality of individual decisions (Heiss et al. 2010, 2013; Abaluck and Gruber 2011, 2013; Kesternich et al. 2013; Ketcham et al. 2012; Kling et al. 2012). Heiss et al. (2009) note the stark growth in Part D premiums for generous coverage and acknowledge that adverse selection may be an important concern in this highly regulated setting. They allow for selection in their simulations of a life-cycle choice model in an environment stylized to reflect the key features of Medicare Part D. Several papers in this literature have suggested that switching costs may be present in the program. In addition to the analysis in Ericson (2013) that documents evidence of insurer pricing strategies consistent with the presence of inertia, Miller and Yeo (2012) assume that the Medicare Part D market exhibits consumer inertia and include a switching cost parameter in their choice model estimated on the market-level data. Further, Abaluck and Gruber (2013) in their analysis of choice inconsistencies allow for switching costs separately from other choice imperfections. Ketcham et al. (2012), on the other hand, suggest that inertia is not of key concern in Part D. Some of the other work analyzing Medicare Part D has looked at the impact of Part D on prescription drug consumption (Yin et al. 2008; Duggan and Scott Morton 2010); the role of low income subsidy regulation (Decarolis, 2013); the welfare of reducing choice (Lucarelli et al., 2012); and the moral hazard response to non-linearities of the contracts (Einav et al., 2013). Duggan, Healy, and Scott Morton (2008) provide an extensive overview of the program's design.

Second, beyond the Medicare Part D setting, the paper is related to the growing empirical literature that analyzes asymmetric information, regulation, contract design, and welfare in both employer-provided and public health insurance settings. Einav, Finkelstein, and Levin (2010a) provide a systematic overview of the literature in this vein. The current paper builds upon the insights in the work on the interaction between adverse selection and minimum standard regulation in Finkelstein (2004), as well as on the interaction between adverse selection and inertia in an employer-provided insurance setting in Handel (2013).

Third, the paper also relates to the growing literature on inertia and defaults in a variety of public finance settings: among recent examples are Chetty et al. (2013) and Beshears et al. (2013), who document such patterns in retirement savings accounts, and Nosal (2012), who estimates switching costs in Medicare Advantage. Finally, methodologically and conceptually, the paper is related to a broad literature in industrial organization that assesses the impact of switching costs and incumbent advantages on market outcomes in a variety of settings (Farrell and Klemperer 2007 provides a survey), as well as the vast literature on the regulation of private markets (see Joskow and Rose 1989; Armstrong and Sappington 2007).

1.2 Institutional Setting and Data

Basics of Medicare Part D

Medicare is a public health insurance program for the elderly and disabled in the U.S. Until 2006, standard Medicare insurance, so-called Parts A and B, covered hospital and physician services, but not prescription drugs. In 2006, Medicare Part D prescription drug coverage was launched as part of the Medicare Modernization Act of 2003, becoming the largest expansion of Medicare since its introduction in 1965. While Medicare bears the greater share of Part D costs (CBO projects 2013 spending on Part D to be $63 billion or 2% of the 2013 budget outlays), the actual administration of the Rx benefit and part of the actuarial risk have been outsourced to private insurers. In 2012, Part D covered around 32 million beneficiaries (Hoadley et al., 2012) - 62% of these were enrolled in stand-alone prescription drug plans (PDPs), which are the focus of the current paper.

Medicare Part D coverage is voluntary and enrollment is not automatic for beneficiaries.[5] Eligible individuals have to actively enroll in one of more than 30 stand-alone Rx plans offered in their state of residence during annual open enrollment period or when they first become eligible, e.g. turn 65. Once enrolled, beneficiaries pay premiums on the order of $\$400 - \$500$ a year, and in return insurers pay for prescription drug purchases subject to a deductible, co-payments or co-insurance, and coverage limits. Beneficiaries stay in their chosen plan for a year and may change their choice during the open enrollment period next year. If beneficiaries make no changes to their plan choice in subsequent years, CMS will continue enrolling them in their first chosen plan unless it is terminated by the insurer. The fact that individuals self-select into plans and have a "default" plan if they do not take any action after their first enrollment will be important for my analysis.

Data – baseline sample

I utilize the detailed administrative data provided by CMS that comprises a 20% random sample of Medicare beneficiaries nationwide for years 2006-2009. The data provides basic demographic and detailed health information about the beneficiaries, the characteristics of all Part D plans available in each region of the country, the enrollment choices of the beneficiaries, and subsequent prescription drug spending for those who enrolled in Part D. I make a number of restrictions to the original sample of 40.3 million beneficiary-year observations to isolate the part of the market where 65 year old and older enrollees self-sort into a cleanly observable set of Part D contracts. These restrictions, which bring the sample down to 5.3 million individual-year observations, are recorded in Table 1.12[6] This baseline sample has individuals that chose to enroll in a stand-alone prescription drug plan and did not receive any additional subsidies from the government that would have distorted the

[5]This is true only for the so-called "regular" beneficiaries, who are the focus of this paper. Beneficiaries eligible for low-income subsidies as well as "dual" eligibles for Medicare and Medicaid are usually assigned to plans automatically by CMS.

[6]I restrict the sample to individuals of age 65 and older residing within 34 Medicare Part D regions or 50 states (Medicare combines some states into the same PDP market), who did not die in the reference year and were originally entitled to Medicare because of old age rather than disability. In other words, I do not include individuals, who may become eligible for Medicare before they turn 65 as part of their SSDI benefit. I further drop observations on individuals that were dual eligible for Medicare and Medicaid in the reference year, since these individuals are assigned to plans by CMS rather than choosing plans on their own. This brings the sample down to 25.6 million beneficiary-year observations. I then eliminate individuals that did not enroll in Part D or were enrolled in Medicare Advantage (or another managed care) option that combines prescription drug coverage with healthcare insurance.

monthly premiums or cost-sharing characteristics of the choices. Since in the econometric choice model I need to observe individual's choices over consecutive years, I also construct a panel sub-sample of the baseline sample. This sub-sample contains individuals whose choices and utilization can be observed from the first year they enter the program to 2009.[7] The panel sub-sample has 3.7 million beneficiary-year observations on approximately 1 million unique individuals.

Table 1.1 provides the summary statistics of the observed demographic and risk related variables for the full sample, the baseline sample and the panel sub-sample. The individuals in the baseline sample are on average 76 years old, 64 % female, predominantly white (95 %), with a risk score of about 0.9 and annual average drug spending of about $1,900. The panel sub-sample looks very similar, albeit a year younger on average and with slightly lower risk scores and annual spending. In comparison to the full sample that includes beneficiaries eligible for both Medicare and Medicaid, non-enrollees, and those who qualified for Medicare before turning 65, the baseline sample has individuals that are older, more often white and more often female. Overall, the baseline sample has somewhat healthier individuals compared to the full Medicare population, as measured by the average risk score of 0.9, which is 10 % below the population average that by construction is 1.

Regulatory environment and the nature of the observed contract space

The nature of Medicare Part D's contract space is driven by a minimum standard regulation. Medicare has designed a so-called Standard Defined Benefit (SDB) for the Part D program and insurers are required to provide coverage that gives at least the same actuarial value as the SDB. The SDB has an unusual design that features a relatively low deductible, flat co-insurance rate of 25% up to the initial coverage limit (ICL) and subsequent "donut hole", or coverage gap, that has a 100% co-insurance until the individual reaches the catastrophic coverage arm of the contract. Figure 1-1 illustrates what these features imply for a beneficiary. Consider an individual who in 2006 was in an SDB contract and purchased prescription drugs for a total of $3,000. Out of this amount, the individual would pay the deductible of $250, then 25% of the next $2,000 up to the initial coverage limit of $2,500, and then 100% of the next $750 in the gap, for a total out of pocket spending of $1,500. The remaining $1,500 would be paid by the plan.

A crucial feature of the institutional setting, which generates variation in contract characteristics, is that insurers are allowed to adjust and/or top up the SDB contract design as long as their contracts cover at least the same share of average spending as the SDB. As a result, contracts offered by Part D insurers are highly multidimensional and vary on a variety of characteristics that differentiate them from the SDB minimum. Some of this differentiation is purely financial - con-

[7]Most differences between the panel sub-sample and the baseline comes from the way CMS draws its 20% random sample of the Medicare population. These samples are only partially based on panel draws and thus not all individuals are observed in every year. For details on the CMS sampling procedures see the Chronic Condition Data Warehouse User Manual v.1.7. Some individuals in the panel sub-sample will be lost if they change from a PDP to a Medicare Advantage prescription drug plan simultaneously with switching from the "traditional" Medicare to the HMO system. Moreover, it is possible that some individuals leave the Part D program altogether; this option is likely to be very rare, however, since these beneficiaries would then face premium penalties if they decide to re-enter the program at a later date. Lastly, some observations will be lost in the panel sub-sample due to individuals dying in years 2007-2009.

tracts can change cost-sharing thresholds, co-pay and co-insurance levels, and may offer coverage in the "donut" hole. Other differentiating features are more related to the quality of the insurance provider.

Despite the multi-dimensionality of contracts and official counts of more than 1,500 contracts in the Part D program, there are three stylized facts about this market that haven't been emphasized in the literature, but allow me to simplify the description of the contract space. First, each insurer in practice offers the same menu of 2-3 contracts in all Part D regions in which it operates. Second, insurers tend to adjust only premiums, but not other contract features across different regions. Third, insurers tend to keep the "types" of contracts in their menu fixed over time, adjusting the key characteristics only to the SDB policy changes.[8] Using these three stylized facts, I classify all contracts into four types. Contracts that offer the standard-defined-benefit level of deductible and initial coverage limit are classified as *Type 1* contracts. *Type 2* contracts offer a reduced deductible (usually reduced all the way to zero), but still the standard level of the initial coverage limit[9]. *Type 3* contracts offer a reduced deductible and partial coverage (usually coverage of generics) in the gap beyond the ICL, while *Type 4* contracts offer a reduced deductible and full coverage in the gap.

In addition to the cross-sectional complexity of the contract space, there is substantial over time variation in the characteristics and premiums of the available plans. This feature of the Medicare Part D program motivates the question, pursued in this paper, about the role of switching costs in determining the allocation of risks across contracts in an environment where the relative attractiveness of choices evolves due to market forces and regulatory interventions. The latter is clearly seen in Medicare's annual adjustments of the SDB cost-sharing thresholds.[10] Figure 1-1 illustrates the annual adjustments in the deductible and the ICL of the standard defined benefit in years 2006-2009. Every year CMS raised the SDB deductible, which made the SDB contract less generous, and at the same time increased the ICL by more than the deductible increase, which raised the actuarial value of the contract. Since beneficiaries with different health risks should be responding differentially to the adjustments in different arms of the standard contract, CMS in theory can leverage the SDB policy not only to change the overall level of coverage, but also to influence the distribution of risks across contracts that alter different arms of the SDB.

Figure 1-2 illustrates the basic evolution of the contract space over time using the 4-type classi-

[8]These empirical observations support the identification of the choice model in Section 1.4 that treats the key contract features as exogenous with respect to individual demand in a given year.

[9]While in theory some of the "type 2" plans would be classified as actuarially equivalent to the minimum standard plan by CMS, this actuarial equivalence holds for average spending. For each individual, however, the zero vs. 295$ deductible means lower spending with a high probability, since many individuals never reach spending above the deductible level. Thus, I consider it plausible to assume that without explicit calculations of the average expected spending, individuals would regard plans that have zero deductible as more generous.

[10]From q1medicare.com: Updating the parameters helps ensure that the government's share of the Part D costs remains constant over time. The annual percentage increase in average per capita Part D spending – used to update the deductible, initial coverage limit, and out-of-pocket threshold for the defined standard benefit for 2010 – is 4.66 percent. The annual percentage increase in the Consumer Price Index – used to update the 2010 maximum co-payments below the out-of-pocket threshold for certain dual eligible enrollees – is approximately 2.65 percent. CMS revised these percentages to correct calculation errors identified following release of the Advance Notice.

fication of plans. Each panel plots the development of enrollment shares, average premiums, as well as average realized drug expenditures. We see that contracts of *Type 2* with a reduced deductible and no gap coverage had the highest enrollment, greater than 60%, in all years. The enrollment share in these contracts increased over time by almost 8 percentage points, mirroring the decrease of 8 percentage points in the enrollment share of the *Type 1* contracts with SDB deductible. The enrollment share in *Type 3* contracts is more volatile over time. Average premiums and average spending have remained relatively stable in *Type 2* contracts, while both grew in *Type 1*, *Type 3*, and especially *Type 4* contracts. The differential growth in the premiums of the different contract types implies that relative premiums changed substantially over time. These regulatory and market-driven trends highlight the dynamic changes in the relative prices and generosity of Part D plans that I will exploit in my empirical analyses.

1.3 Descriptive evidence

1.3.1 Adverse selection in Medicare Part D

As is typical for insurance markets, Medicare Part D plan providers cannot price-discriminate individuals according to their risk profiles. The consequence is that while all enrollees in a given plan pay the same premium, they impose different costs on the insurers. Moreover, the more comprehensive insurance contract the insurer offers, the more likely it is to attract an adversely selected risk pool of costlier individuals, since healthier individuals will not be willing to pay high premiums for coverage they don't expect to need and will not want to pool with costly beneficiaries. In this section, I document the presence of adverse selection in Medicare Part D in several steps. To begin with, I present the cross-sectional correlation test as described in Chiappori and Salanie (2000) using the ex-post realized drug expenditures. This exercise detects the presence of asymmetric information and follows the well-established testing literature reviewed in Einav et al. (2010a). I then provide two pieces of additional evidence that help disentangle selection from moral hazard, which is a common concern in health insurance. First, I repeat the Chiappori and Salanie (2000) test using ex-ante information about individuals' health summarized in a risk score index. Second, I document the presence of two selection death spirals in the early years of the program.

Figure 1-4 illustrates the first positive correlation test graphically. It plots the average realized drug spending in each region by contract type in years 2006 and 2009. We can see the stark differences in the expenditures in the more and less generous contracts. Especially striking is the panel for 2006, as it contains data for *Type 4* contracts that have realized average annual spending on the order of $\$4,000$ in all regions, as compared to $\$1,500$ in the *Type 1* contracts. The differences in expenditures appear to shrink slightly in 2009 as compared to 2006, with the high risks of the former *Type 4* contracts now integrated into the rest of the market.

Table 1.2 presents the formal specification of the test for the presence of asymmetric information. The test is done using the baseline sample pooled for years 2007-2009 that includes observations for which risk scores and claims are accurately measured for the whole year. Since insurers are allowed

to price the same contracts differently in different regions, the test controls for region fixed effects. The regression specification takes the following form:

$$Y_{irt} = \alpha_r + \delta_t + \sum_{k=2}^{k=4} \beta_k \mathbf{1}\{ContractType_{it} = k\} + \epsilon_{irt} \tag{1.1}$$

where i indexes individuals, r indexes regions and t indexes years. I use realized total drug spending as the first outcome variable in Column (1). Since the goal of the exercise is to test whether higher risk individuals sort into more generous contracts, the spending variable does not account for the cost-sharing provisions of the plans. The results suggest that more generous contracts have higher spenders. For instance, contracts with full gap coverage attract individuals with realized drug spending that is more than a standard deviation higher than in the plans with minimum standard coverage.

Since the correlation tests that use the realized spending as the outcome variable capture both adverse selection and moral hazard, I repeat the testing exercise using risk scores as the outcome variable in Column (2). Risk scores are constructed using ex-ante diagnostic information from Medicare Part A/B (hospital and physician) claims and therefore do not reflect any effects of plan structure on spending that may show up in the tests that use drug spending as the outcome variable. Using the risk score measure gives qualitatively similar results, although the magnitude of differences is smaller. To give a more meaningful interpretation to the results with risk scores, in Column (3) I project drug spending onto risk scores and use the projection as the outcome variable in the correlation test. This exercise expresses the risk scores in dollars of expected spending. Since risk scores do not have the moral hazard aspect, this projection gives me the *lower bound* on how much of the estimated differences between the realized spending in different types of plans can be attributed to adverse selection on ex ante observed (and thus potentially priceable) risk. The results remain qualitatively similar. Although this measure accounts for only about a third of the differences that were observed in the first regression, the differences are still large in the absolute sense, corresponding to more than 100% of the average annual premiums.

Figure 1-5 illustrates the moral hazard-free results graphically. Instead of comparing just the average expected risk, this figure plots the whole empirical CDF of the ex-ante risks in different types of contracts in years 2006 and 2009. We can clearly see that, indeed, the whole distribution of risks in the more generous contracts is shifted towards having more mass of higher risks. The difference between the 2006 and 2009 panels demonstrates how the high risks that selected into *Type 4* plans in 2006 got re-incorporated into the system after these plans no longer provided full gap coverage. I will discuss the role that switching costs may have played in this process in the counterfactual analysis in Section 1.5.2.

In addition to this evidence of cross-sectional adverse selection, Medicare Part D illustrates rarely observed evidence of dynamic adverse selection. Despite the extensive efforts of the regulator to incentivize private insurers to offer Part D contracts with full gap coverage, such contracts were discontinued after 2007. The selection spiral happened twice, since different insurers attempted

offering full gap coverage in 2006 and 2007, both of them discontinuing full gap coverage after one year of operation. Figure 1-2 documented that *Type 4* plans with full gap coverage experienced annual claims that were about twice as high as the claims in plans without gap coverage. Subsequently, their premiums nearly doubled and enrollment dropped. Up until today no plans in Medicare Part D program offer full coverage in the gap.[11] The Medicare Part D environment thus illustrates a rare setting, where an off-equilibrium plan was in practice offered on the market and then rapidly unraveled.

1.3.2 Switching costs in Medicare Part D

One objective of this paper is to analyze how risk-sorting among contract types documented above changes in response to regulatory and market-driven adjustments in the contract menu. We would expect changes in the relative prices or generosity of the available insurance contracts over time to induce a re-allocation of risks among them. In the presence of high switching costs, however, such re-allocation may be completely muted. In this section I begin exploring this interaction mechanism by first documenting evidence for the presence of switching costs in Medicare Part D.

In general, documenting evidence of switching costs is challenging, since we need to distinguish between the "structural" state dependence and unobserved persistent individual heterogeneity (Heckman, 1981; Heckman and Singer, 1986; Heckman, 1991; Honoré, 2002; Honoré and Tamer, 2006; Dube et al., 2010). Several features of the Medicare Part D environment, however, render themselves well to such analysis. First, due to the regulatory changes, entry and exit of plans, as well as substantial market-driven over time adjustment in prices and characteristics of contracts (some of it strategic, see Decarolis (2013)), there is pronounced non-stationarity in the observed contract space of the program. Second, the Part D environment allows observing all individuals making their first choice in 2006, as well as younger individuals making their first choices in years 2007-2009 from the adjusted contract menus. In this section, I utilize these features of the Part D environment to provide descriptive evidence of choice behavior consistent with the presence of switching costs as separate from persistent individual heterogeneity.[12] This approach to documenting switching costs is in the spirit of Handel (2013)'s approach to documenting inertia in employer-provided health insurance.

I find four descriptive patterns in the data consistent with the presence of significant switching costs in Part D. First, in Table 1.3, I document that in each year of the program about 90% of

[11] In addition to a strongly selected enrollment pool, offering full coverage in the gap by regulatory construction delayed the start of the catastrophic coverage arm of the benefit, in which the insurers are only responsible for 15 % of the costs. Since the catastrophic coverage limit starts once a certain out-of-pocket spending amount is reached, offering coverage in the gap significantly slowed down the accumulation of the out-of-pocket amount. The government realized that this created strong disincentives for insurers to offer coverage in the gap and conducted an experimental "demonstration" that changed the structure of the federal reinsurance for plans that offered gap coverage. However, these changes in incentives were not sufficient to overcome strong adverse selection into these plans.

[12] Section 1.4.1 then presents a more formal discussion of the issues related to the identification of the state dependence parameter in dynamic discrete-choice panel data models.

individuals enroll in their "default" plans; for individuals whose default plans significantly change their financial characteristics (and thus their type) this probability is still around 80%. Since premiums and contract characteristics change substantially from year to year even if plans do not change their "type", the high persistence in choices suggests that switching costs may be present; this evidence alone, however, could just point to very persistent preferences.

As the second piece of evidence, I compare the choices of the newly entering and existing enrollees in different years. The results are recorded in Table 1.4. I focus on the individuals that can be tracked from their first entry continuously to 2009 and whose default plans' types were not changed by insurers throughout the observed years. This isolates individuals whose choices are not conflated with substantial supply-induced re-classification of plan types. Two patterns are pronounced in the data and consistent with the idea that switching costs play an important role. First, enrollment shares over time for a given cohort tend to be closely related to the choices and market conditions of the first year in which the cohort entered the program. Second, the choices of different cohorts in the same year are very different. For example, comparing the choices in 2008 of the cohort that entered in 2006 and the cohort that entered in 2008, we see that the 2008 cohort is almost twice less likely to enroll in the least generous *Type 1* plan than the 2006 cohort in 2008: 10% vs 19 % enrollment share.[13] Another persistent difference in choices is visible for the 2007 cohort, which in 2007 was much more likely to select the *Type 3* plan with partial gap coverage than cohorts entering in other years. The inertia in choices is certainly not complete, as for all cohorts we do observe the tendency of enrollment to move in the direction of choices made by the newly entering cohort. The adjustment, however, is much slower than if the individuals with default plans behaved as if they were choosing anew.[14]

Comparing the choices of different cohorts by their choice of different insurer brands rather than contract types paints a similar picture. Figure 1-6 records the enrollment shares of the two biggest insurers in the sample for each year 2006-2009. The enrollment shares are shown separately for the 65 year olds, who are entering the program anew, and older enrollees with incumbent plans. We see a striking difference in the 2009 choices. In this year, one of the insurers ("Insurer B") lost almost all of its market share with the new enrollees. Only about 5% of the 65 year olds chose to enroll in the plans offered by this insurer. Among the continuing cohorts, its enrollment share also fell, but not nearly as dramatically. It remained higher than 20%, implying that in 2009 the existing cohorts were four times more likely to be enrolled with Insurer B than the new enrollees.

In addition to the descriptive evidence on switching costs in Table 1.3, Table 1.4, and Figure 1-6, I use a simple conditional logit regression to test whether there are statistically significant

[13]Note that this pattern undermines an argument that individuals in the existing cohorts are choosing differently, because they are slightly older and thus have higher drug spending risk, not because they have switching costs. In this particular example, older cohorts are choosing *less* generous coverage more often.

[14]Note that we can see evidence of switching costs in this type of comparisons only if the relative attractiveness of choices changes over time. When this is not the case, as in years 2008-2009, where the structure of contracts and premium levels were similar, we see very little difference in the 2009 choices of the cohort that entered in 2009 and a year earlier in 2008.

differences in the price sensitivity of the new and continuing enrollees.[15] Under the null hypothesis of no switching costs, we would expect the coefficients on plan premiums for new (65 y.o.) and existing enrollees of similar age (66 - 70 y.o.) in the same year to be very close to each other. The estimates presented in Table 1.5, however, allow me to reject this null. The coefficient on premiums in the utility function, which is proportional to the elasticity of enrollment probability, is higher in magnitude for 65 year olds than for 66-70 year olds in all model specifications in years 2007-2009 of the program. This does not hold in 2006 when beneficiaries of all ages are entering the program anew. Furthermore, the estimates of the price coefficient are virtually identical for each age group among 66-70 year olds. The latter finding suggests that the difference between the estimated price sensitivity for the new and existing enrollees is not driven by age differences per se, but instead are related to the lack of switching costs for the 65 year old beneficiaries.

1.4 Empirical model

1.4.1 Specification and identification of the contract choice model

The descriptive evidence in Section 1.3 has documented that adverse selection and switching costs are present in Medicare Part D. This evidence alone, however, doesn't allow quantifying the economic significance of these market imperfections, or the extent of their interaction with regulatory policies. To quantify the importance of information about risk for contract choices, as well as how much the minimum standard regulation affects the allocation of risks in the presence of switching costs, I formulate an econometric model of how individuals choose which contract to enroll in. The model specifies the choice decision as a function of the information about the individual health risk, the switching cost, and heterogeneity in individuals' preferences for different features of the contracts. The model takes a contract-valuation approach, allowing me to make relatively few assumptions on how individuals interpret the financial features of the contracts and on the individuals' information set about risk at the time of choice.[16] This comes at the cost of not recovering deeper utility primitives, such as risk aversion. Revealed valuation of contract characteristics, however, is sufficient for assessing how choices and risk allocation would change with costless switching and in response to policy-driven changes in contract characteristics that are reasonably within the realm of the observed contracts. These are exactly the type of counterfactuals that I consider in Sections 1.5.2 and 1.5.3.

Specification Each year t an individual i who lives in region r and is enrolled in the Medicare Part D stand-alone prescription drug program chooses among J_r plans offered by B insurers. Each insurer b typically offers a menu of up to three plans of different types in each region. The plans of the same type offered by the same insurer are likely to have the same characteristics, but different

[15]The exact specification is discussed in the notes of Table 1.5 that reports the results of this regression.

[16]Both the contract-valuation and the realized utility approaches to modeling of health insurance choices are common in the literature. Einav et al. (2010a) discuss the trade-offs between these approaches, illustrating the contract-valuation approach on the work of Bundorf et al. (2012) on pricing in the employer-provided health insurance.

premiums in different regions. Some plans are offered only in a subset of regions. These two features of the program imply that J_r varies by region. The plans in each J_r can be projected into the same set of observed characteristics. While in principle the plans could be characterized by a very high-dimensional vector of characteristics available from the administrative records, in practice I have to take into account which characteristics of the plans are feasibly observed by beneficiaries when they are making their choice. I let individual i's utility from choosing plan j (where "plan" is region-specific, so r-indexing is suppressed) in year t be given by:

$$u_{ijt} = -\alpha p_{jt} + \beta_{it}\phi_{jt} + \gamma_{it}\mathbf{1}\{\text{Default}\}_{ijt} + \epsilon_{ijt} \tag{1.2}$$

$$\epsilon_{ijt} \sim \text{iid Type 1 EV}$$

Utility thus depends on the annual premiums charged by the plan in a given region in a given year p_{jt}, the characteristics of the plan ϕ_{jt}, and whether j was a default plan for individual i in year t, where the default plan is usually the plan chosen in $t-1$.[17] Individuals are assumed to choose a plan that gives them the highest utility. An important assumption that is implicit in this formulation of the utility function is that individuals are myopic in their choice of plans. In other words, I am assuming away a possibility that individuals, for example, expect a plan that is cheap today to become very expensive tomorrow and thus choose a suboptimal plan today to avoid paying the switching cost tomorrow. Further, this formulation assumes that individuals choose the option with the highest "perceived" utility, which may not necessarily correspond to the highest "objective" valuation of plans as financial contracts (indeed, Abaluck and Gruber (2011, 2013) suggest that beneficiaries are choosing their plans inconsistently with the objective efficiency frontier). For the analysis of risk-allocation and choices in counterfactual scenarios, however, this "subjective" utility is exactly the object of interest.

I let the characteristics component of the utility function ϕ_{jt} include the following plan features:

$$\begin{aligned}\beta_{it}\phi_{jt} &= \beta_{1it}\text{Deductible}_{jt} + \beta_{2it}\text{ICL}_{jt} + \beta_{3it}\mathbf{1}\{\text{Partial coverage in gap}\}_{jt} + \\ &+ \beta_{4it}\mathbf{1}\{\text{Tiered Cost Sharing}\}_{jt} + \beta_{5it}\mathbf{1}\{\text{LIS eligible plan}\}_{jt} + \beta_{6i}\mathbf{1}\{\text{Brand}\}_{j}\end{aligned} \tag{1.3}$$

This specification assumes that conditional on insurer, differences in plans can be accounted for by the deductible level, the initial coverage limit and a set of indicators on whether or not the plan offers partial coverage in the gap, whether the plan uses fixed dollar co-payments or co-insurance percentage, and whether the plan is eligible to enroll individuals with low-income subsidy. The included characteristics capture a substantial amount of variation among plans, since many features that are not explicitly included, such as the quality of services, pharmacy network quality,

[17] To construct the "default" variable I use the administrative records of which plans were renewed over time, which plans were discontinued and which plans were consolidated. In cases where plans were renewed or consolidated, CMS would default individuals in the same (if renewed) or designated new (if consolidated) plan if individuals took no action to change their choices. In rare cases when plans were terminated, individuals are recorded to have no default plan in that year.

and the generosity of drug formularies are insurer-level rather than plan-level characteristics. The features of the plans included in this specification correspond to the information that individuals had readily available from front-end consumer advertisement materials (e.g. fliers distributed at retail pharmacies) by a typical insurer on its 1-3 different contracts.

Beneficiaries in health insurance markets differ in two key ways - in their preferences for different contract features, as well as in how costly they are for the insurers. Individual preferences, in turn, may reflect both individual health risk as well as horizontal tastes and risk aversion that may or may not be correlated with the expected drug spending. To capture these features of insurance demand in the model, I first allow for rich observed heterogeneity of preferences in the specification of marginal utility from contract features. Importantly, I allow individual preferences to depend on the individual's health risk. I use the full set of demographics observed in the data - age, gender, and race; as well as proxies for expected spending - risk scores and an additional flag for having end-stage renal disease diagnosis, which identifies especially high risks. Vector D_{it} records this demographic information and risk measures: $D_{it} = \{\text{age}_{it}, \text{gender}_i, \text{race}_i, \text{risk score}_{it}, \text{esrd indicator}_{it}\}$. The reduced-form evidence for adverse selection in Section 1.3.1 suggest that individuals may have more private information about their expected spending than what is accounted for by risk scores. Thus, I additionally allow for unobserved heterogeneity in preferences for the key financial features of contracts. This is achieved by specifying random coefficients on three key features of the contracts: deductible, initial coverage limit, and partial gap coverage. The distribution of the random coefficients is assumed to be normal. I interpret this unobserved heterogeneity as likely stemming from the private information about health risk not captured in risk scores as well as from the heterogeneity in risk aversion.

The assumption behind this specification of heterogeneity is that individuals are aware of their previous medical diagnoses and what these diagnoses typically imply for drug expenditures. Note that this specification doesn't impose that individuals have precise expectations about their ex-post realized spending (the spending is not used in the estimation of the model), but rather have a general understanding that certain diagnostic groups on average cause higher drug spending. An additional dimension of heterogeneity is conceivable with respect to the switching cost. If switching costs are, for instance, interpreted as search costs, we may think that, for example, older and sicker individuals have higher switching costs. To capture this possibility, I allow the switching costs to differ among individuals of different demographic groups and risk types. All in all, the coefficients on the contract characteristics and the lagged dependent variable are specified as follows:

$$\beta_{it} = \pi^{\beta} D_{it} + \psi_i^{\beta}, \text{ where } \psi_i^{\beta} \sim N(\psi^{\beta}, \sigma^2) \tag{1.4}$$

$$\gamma_{it} = \pi^{\gamma} D_{it} + \psi^{\gamma} \tag{1.5}$$

Assuming that an individual chooses the plan that maximizes his or her utility, the model allows expressing the probability of the beneficiary choosing different plans in his or her choice set as a function of parameters. We can then use, for example, the maximum likelihood estimation approach

to find the values of the parameters that best rationalize observed choices.[18] Since in the data I can track the same individuals making several consecutive choices, the estimation utilizes this panel structure, explicitly modeling the probability of a sequence of choices. While assuming the extreme value distribution for the taste shocks produces a closed-form probability expression conditional on the realization of the random coefficients, the unconditional probability involves integrating out the normally-distributed random coefficients. The latter implies that there is no analytic closed-form solution for the probability integral that is part of the log-likelihood function. Thus, the model is estimated using a simulated maximum likelihood (MSL) procedure as described in Train (2003, 2009) and Hole (2007).

Identification. The identification of the parameters relies on several fairly unique features of the data. First, to recover the switching costs parameter γ, we have to consider two issues: distinguishing between the "spurious" versus "structural" state-dependence and the initial conditions problem.[19] The inclusion of the unobserved individual heterogeneity through random coefficients into the model addresses the first issue in a way that is standard in the literature. The assumption is that the normal distribution of the random coefficients correctly captures the heterogeneity, and thus the lagged dependent variable parameter estimates the "structural" part of state-dependence. Moreover, in 2007, 2008, and 2009, there were cohorts of 65 year olds that first became eligible for Medicare and entered the Part D program anew without switching costs. This implies that in years 2007-2009 of the data I observe individuals choosing with and without switching costs from the same menu of contracts. The latter feature greatly aids in separating the persistent individual heterogeneity from the switching friction. The initial conditions problem does not arise in the current setting, as I observe the first choices for all individuals in the estimation sample, since the year of the program's launch - 2006 - is recorded in the data. Finally, the descriptive evidence in Sections 1.2 and 1.3 suggests that there is substantial variation in the prices and the characteristics of plans in each year of the program. Such variation is important, since if the environment were very stable, we couldn't expect to observe any changes in choices either with or without costly switching.[20] These features of the data combined, allow for the identification of preferences separately from the switching costs that are hard to distinguish in more typical observational choice data settings.

[18] As these derivations are standard for a mixed logit model, I omit the details in the paper.

[19] The concern in the first issue is that the lagged dependent variable in the utility function, which is capturing the switching cost, will be correlated with (or rather directly a function of) an individual-specific preference parameter. To illustrate, in a generic binary non-linear dynamic panel model, this would imply that in $y_{it} = 1\{\beta x_{it} + y_{i,t-1}\gamma + \alpha_i + \epsilon_{it} > 0\}$, $y_{i,t-1}$ is a function of α_i. Thus, if α_i is unaccounted for and left in the unobserved part of the utility function, the identification assumptions about ϵ_i will be violated. The literature on the non-linear dynamic panel data discusses the two broad approaches to this problem - assuming a parametric distribution for the unobserved individual heterogeneity ("random effects"), or trying to difference out the individual effects without functional form assumptions ("fixed effects"), where the latter approach encounters a lot of challenges given the non-linear nature of the model. Honoré (2002) and Honoré and Tamer (2006) provide an excellent discussion.

[20] Dube et al. (2010) discuss the importance of observing variation in the choice set for the identification of structural state-dependence; they utilize promotions as generating such variation.

To recover the preferences of different demographic and risk groups for contract characteristics, I take advantage of the substantial variation in the data stemming from several sources. First, there is rich observed heterogeneity in individual demographics and health risks across the country. Second, there is variation in the characteristics and prices of contracts offered within a given region in a given year: each individual faces more than 30 contracts offered by multiple insurers. Further, there is substantial variation in the choice sets across different geographic regions, and within regions over time. The cross-sectional variation in the non-price features offered by different plans (such as zero deductible) is generated by the insurers' strategy of offering menus of several vertically-differentiated contracts. This strategy appears to be empirically stable over time, suggesting no contemporaneous responses to aggregate demand shocks. A lot of time-series variation in non-price contract features is generated by the changes in the minimum standard policy that annually adjusts the standard deductible and initial coverage limits. The variation in premiums stems from two sources. First, insurers set different relative prices for the two or three contract types in their contract menus. Second, insurers set different prices for the same contracts in different geographic regions.

This rich variation in premiums comes from observational data and not pricing experiments, suggesting that endogeneity concerns are warranted. We may first be concerned about the endogeneity of the insurers' pricing decisions to the aggregate of beneficiary choices. The aggregate premium endogeneity concerns are partially ameliorated by two observations. First, the premiums faced by individuals reflect only a small fraction of the actual prices charged by the insurance providers to Medicare. These premiums are constructed through a bidding mechanism that ties the individual premiums to the average of prices charged by all insurers to Medicare; thus, insurers do not know exactly which premiums individuals will face for the plans in advance of setting prices. Second, Medicare imposes guidelines as to what type of costs insurers are allowed to include in the calculation of their premiums and how these costs should be related to the insurers' risk pool.

In addition to aggregate endogeneity, the stochastic component of the utility function ϵ_{ij} may include unobserved characteristics of contracts that are correlated both with premiums and individual choices, leading to an omitted variables bias. For example, an insurer could be advertising a particular contract in its menu more than other contracts and setting the price of this contract higher/lower because of that. Note that conditioning on the insurer-specific fixed effects does not resolve this issue if insurers are advertising a particular contract in their menu. One example of such situation in the Medicare Part D setting is the endorsement of selected contracts by a well-known third party. Several contracts in the portfolio of one insurer were endorsed by the American Association of Retired Persons; as the names of the contracts and insurers are not observed in the administrative data due to commercial privacy restrictions, we cannot include a variable measuring this endorsement directly. At the same time, it would be natural to assume that the AARP endorsement both leads individuals to select this contract with a higher probability and allows the insurer to raise prices either to exploit the less elastic demand, or to cover costs that may arise from the marketing relationship with the AARP. Having such unobserved characteristics in the

stochastic portion of the random utility specification violates the assumption of no correlation between the observed and unobserved components of utility. The standard approach in the Medicare Part D literature (Abaluck and Gruber (2011); Heiss et al. (2013)) has been to assume that the rich observed characteristics capture all the relevant information about choices. In this paper I utilize an instrumental variables strategy based on the observations of contracts' expenditures in the administrative data to improve upon this approach.

To correct for potential endogeneity, we need an instrumental variable that affects the contract premiums, but is not correlated with the contract characteristics not observed in the utility function. Since the costs of insurance contracts depend almost entirely on the prescription drug claims submitted by their enrollees, it is natural to consider these claims as a potential source of the variation necessary for the IV strategy. In particular, as the regulatory guidelines suggest, and the results of the pricing model in Section 1.4.2 confirm empirically, contract prices are strongly correlated with the lagged mean realized claims in the contract even after we condition on its key financial characteristics. To utilize the lagged mean realized claims in each contract as an instrumental variable, we need to assume that the variation in this variable is independent of the unobserved contract characteristics conditional on the observed contract characteristics. For example, we need to assume that the AARP endorsement does not affect the level of realized risks in the contract. This assumption appears plausible, as long as we believe that the unobserved characteristics do not screen risks. The latter is indeed very likely, as the reduced-form analysis of Section 1.3.1 suggests that risk-screening happens primarily on the gap coverage margin.

To operationalize the instrumental variables estimation, I utilize the control function approach (Petrin and Train, 2009). Formally, the premium for contract j is a function of observed contract characteristics ϕ_j, a variable that affects premiums, but doesn't otherwise affect the choice decisions z_j , and the remaining unobserved term κ_j (Section 1.4.2 discusses the details of the pricing function in the Part D setting):

$$p_j = f(\phi_j, z_j, \kappa_j)$$

The endogeneity concern arises if κ_j is correlated with ϵ_{ij} in the utility function. Assuming linearity and additive separability of the unobserved component we have:

$$p_j = \lambda_\phi \phi_j + \lambda_z z_j + \kappa_j$$

which is the first stage familiar from a linear IV model. As the choice model is not linear in price, however, the 2SLS technique cannot be applied. One alternative is the control function approach.[21] The idea of this approach is to empirically estimate κ_j and condition on it (or its function) explicitly in the utility function. In practice, κ_j is calculated as the residuals of the first-stage regression of premiums on the observed contract characteristics included in the utility function and the lagged

[21]The other alternative would be to use the Berry et al. (1995) method. This method is somewhat problematic in the current setting, as a large number of contracts has a very small market share. Aggregating the "fringe" contracts, however, would be counterintuitive for the switching cost analysis, as the small market share for some of these contracts stems precisely from the inertia friction.

claims instrument.[22] In the second step, a linear control function $CF = \omega\hat{\kappa_j}$ is included into the utility function:

$$u_{ijt} = -\alpha p_{jt} + \beta_{it}\phi_{jt} + \gamma_{it}\mathbf{1}\{\text{Default}\}_{ijt} + \omega\hat{\kappa_j} + \epsilon'_{ijt}$$

I assume that the stochastic part of the utility function ϵ'_{ijt} has an iid extreme value component that is idependent of κ_j with the remaining component distributed jointly normal with κ_j. This assumption returns a mixed logit model with mixing over the selected characteristics of the contract as well as the error component (Villas-Boas and Winer, 1999).

1.4.2 Empirical model of contract pricing

To analyze the effects of switching costs and the minimum standard regulation on risk-sorting in Medicare Part D in Sections 1.5.2 and 1.5.3, I need to account for how insurers may adjust the contracts they offer in response to the hypothesized changes in the environment. One standard approach at the supply-side analysis would be to assume that premiums are determined as an outcome of a pricing game, such as Bertrand, and are set as a mark-up over the marginal cost. Under a given assumption about the pricing game, the level of mark-ups can then be recovered from observed prices and the estimated elasticities of demand. Making an assumption about the type of game that insurers play and deriving mark-ups using this standard methodology is very problematic in the Part D setting due to the substantial regulatory intervention into insurers' price-setting. The premiums that individuals face on this market are not set directly, but are the outcomes of a "bidding" mechanism run by Medicare. This mechanism determines the payments that insurers get from the enrollees in premiums and from the government in subsidies.[23] Building a pricing model that accounts for this complex regulatory mechanism in Medicare Part D is beyond the scope of this paper.[24]

Instead, here I utilize a stylized approach to empirically relate the premiums faced by the beneficiaries to the risk portfolio and the characteristics of plans.[25] While this approach does not impose an explicit model of competition on the premiums, it does rely on the information about how the regulator outlines the pricing process for the plans. When insurance plans submit their annual bids to the Part D program, Medicare requires them to "justify" the economic validity of the bids. Participating insurers have to fill out worksheets that provide information about the spending experienced by the current enrollees in a given plan in the previous year and how the

[22] In cases where no lagged claims are available, I use average claims in the region as the instrumental variable. The average regional claims proxy for the insurer's expectations of the contract's costs.

[23] At the same time, the bidding mechanism determines which plans are eligible to enroll low-income individuals. Decarolis (2013) shows that the latter aspect may significantly distort pricing incentives of insurers.

[24] We pursue the analysis of incentives and the welfare loss created by the pricing and subsidy regulation in Medicare Part D in Decarolis, Polyakova, Ryan (2013)

[25] As is again standard in supply-side analysis, I treat the non-price characteristics of the contracts as given. To the extent that insurers in Part D offer a fairly stable menu of contracts that follow the regulatory changes in the minimum standard over time, the assumption of exogenous key financial characteristics appears reasonable for the short-run counterfactual analyses.

plan projects these spending will change in light of any planned changes in plan characteristics (usually those driven by changes in the minimum standard regulation). I therefore include the moments of the lagged spending distribution and the key financial characteristics of the plans as the primary components of this hedonic-style pricing regression. Medicare allows plans to include administrative costs and desired profit margins for the plans, which I assume are insurer-specific and so can be picked up by insurer fixed effects. The full specification takes the following form:

$$\begin{aligned} E[Y_{jbt}|.] &= \alpha_b + \delta_r + M'_{jbt-1}\beta + \\ &+ \gamma_1 Ded_{jbt} + \gamma_2 ICL_{jbt} + \gamma_3 1\{PartialGap\}_{jbt} \end{aligned} \tag{1.6}$$

where j indexes plans (where "plan" is region-specific), b indexes insurers (brands), r indexes 34 Part D regions, t indexes years. Y_{jbt} is the annual premium charged to the beneficiaries by plan j of insurer b in year t. Vector M contains several moments of the distribution of drug expenditures experienced by plan j in year $t-1$, including the mean, the standard deviation, the inter-quartile range and the tail percentiles. As Table 1.13 shows, the moments of the lagged risk distribution and the key characteristics of the plans together with the region and insurer fixed effects account for 80% of the variation in the data on premiums over years 2007-2009. As expected, plans with higher lagged mean annual spending, or plans that offer more coverage through lower deductible, higher ICL or partial coverage in the gap are more expensive.

1.5 Results

1.5.1 Parameter estimates and assessment of fit

Parameter estimates Table 1.6 compares parameter estimates from several specifications of the choice model. We observe that instrumental variables strategy increases the magnitude of the price coefficient, as we would expect if the unobserved characteristics in the utility function are positively correlated with the premium and the choice probability. The coefficient increases in magnitude by 12%, suggesting that the observed characteristics of the contracts capture most of the choice-relevant information about contracts and that the endogeneity concerns are limited. Table 1.7 records the detailed parameter estimates of the preferred model specification with the control function instrumental variables approach and an extended number of insurer-specific constants. The willingness-to-pay estimates in the following discussion are calculated by dividing the parameter estimate for a given plan characteristic or demographic interaction by the coefficient on price. The first estimate of interest is the magnitude and the heterogeneity of the switching cost. The switching cost is estimated to be large, but not to vary among demographic and risk groups in an economically significant way. For example, a 75 year old white female with no end stage renal disease and an average health risk (i.e. risk score equal 1) with expected total drug spending of about $\$2,190$ is estimated to face a switching cost of $\$1,164$; while an 80 year old white male with

no ESRD and twice the expected risk is estimated to face a switching cost of $1,253.[26] While there is some risk-related heterogeneity in the magnitude of the switching friction, it is not large enough to suggest that whether lower or higher risks tend to stay in their plans could drive the selection patterns. I estimate that switching cost are increasing in the observed risk at the rate of $74 for an additional unit of risk score; in other words, an individual with twice as high expenditure risk has a 6% higher switching cost. The switching cost is also increasing with age at the rate of about $3 per year. Both of these appear intuitive, since we would imagine that it is somewhat costlier to switch plans for older and sicker individuals.[27]

The second set of estimates sheds light on the role of information about risk for the choice of contracts. Consistent with the reduced-form tests for adverse selection in Section 1.3, I find that beneficiaries with higher health risk value the generosity of coverage more than individuals with lower risks. For instance, individuals with a risk score that is twice the Medicare average, are willing to pay $22 more, on average, for each additional $100 of the initial coverage limit, about $120 more to enroll in plans that have fixed co-pays rather than co-insurance, and about $230 more, on average, to be in a plan that offers partial gap coverage, than otherwise observationally identical beneficiaries with average risk.[28] Moreover, consistent with the hypothesis that beneficiaries have private information about their expected spending beyond the diagnostic information accounted for by the risk scores, I estimate a large degree of heterogeneity in the valuation of the key financial features of the contracts. For instance, the estimate of the standard deviation in the valuation of the partial gap coverage is $250 dollars. This implies that even for individuals within the same demographic and risk group, the differences in the valuation of the coverage in the gap may be very large. Put differently, even if my diagnostic information does not suggest that I am a particularly high risk individual, I may have private information that my expected spending is going to fall into the donut hole, or I may be very risk averse, warranting me to pay around 50% more in premiums to have a contract with partial coverage in the gap.

Model fit and descriptive patterns generated in the model To assess whether the model

[26]The estimated order of magnitude is roughly similar to other findings in the health insurance literature. Thus, Handel (2013) estimates the switching costs to be about $2,000 in the context of employer-provided health insurance; Nosal (2012) estimates the switching cost in Medicare Advantage health plans for seniors to be about $4,000. Using a different choice model, Abaluck and Gruber (2013) estimate the switching costs in Medicare Part D to be on the order of $600 – $700, which is lower than what my estimates suggest in the willingness-to-pay terms, although they similarly find that beneficiaries are roughly 500% more likely to choose a plan that is the "default" plan for them in a given year. Using aggregate data and a dynamic demand model, Miller and Yeo (2012) estimate the switching cost in Part D to be $1,700, which is higher than my estimates.

[27]This intuition holds, however, only if individuals choose plans themselves. We can also imagine that for the very sick, it is their relatives who are choosing the plans and so the switching costs may in fact appear lower. Although the coefficient is not estimated precisely, such logic may apply for individuals with ESRD, whose switching costs are estimated to be $130 lower.

[28]Interestingly that conditional on risk, I find that older beneficiaries are willing to pay more for plans with higher initial coverage limit and coverage in the gap, although they are less averse to having a higher deductible. The former aspect is consistent with findings in an (overall much younger) US employer-provided setting in Geruso (2013), where older individuals appear to prefer more comprehensive plans conditional on their expected risks.

can capture the key patterns in the data, for each beneficiary-year observation I simulate contract choices using the estimated parameters and the observed information about each beneficiary-year. Figure 1-7 illustrates the predicted and the observed enrollment shares for a set of 90 prescription drug plans with the highest enrollment in the sample.[29] Visually, the model performs well in capturing individual preferences for specific plans. Based on the simulated contract choices, I can also calculate the average expected risk that is predicted by the model for each plan. Figure 1-7 illustrates the result of this exercise. The very close fit of the observed and predicted expected risk for each of the 90 plotted plans suggests that the model performs well in capturing the role of risk in the selection of plans.

Table 1.8 summarizes the fit of the choice model on the pooled in-sample data. I report the simulated and the observed values for three moments in the data - enrollment shares, average ex-post spending, and average risk scores along the 4-type plan typology and for top 2 insurer brands. The model fits these moments of the data very well. Enrollment shares and risk scores that are used as inputs in the simulated maximum likelihood procedure fit closely to the data in this pooled representation. This is encouraging, as it indicates that the model is able to capture the key patterns of interest in the data. Since the counterfactual simulations in the next two sections rely on the ability of the model to capture the role of switching costs and the information about expected spending for individual choices, I perform two more checks of the model that test these two aspects. In particular, I generate the key descriptive patterns from Section 1.3 in the model. First, I check whether the model accurately predicts that individuals are likely to choose their default contracts. Table 1.9 records the fraction of beneficiaries in-sample that are predicted to choose their default contracts. The simulated default choice share is about 4 percentage points lower than the observed share, e.g. 86% predicted vs. 90% observed share of default choices in 2007. The difference appears to not be economically meaningful, as the model still predicts that an overwhelming share of beneficiaries, more than 85%, choose default plans every year.

Second, I check whether the model is able to predict the differences in the distribution of risks across contracts of different generosity. This tests whether the model can generate the same adverse selection patterns as the ones observed in the data. Figure 1-8 plots the simulated empirical CDFs of risk by the type of contract for each year separately. The graphs clearly show the adverse selection patterns, with more generous contracts enrolling higher expected risks. The distributions predicted by the model's simulated enrollment choices look quantitatively similar to the distributions of risk generated in the observed data (plotted in Figure 1-5). In the counterfactual simulations in the next two sections, I will analyze how these distributions change in response to the reductions in the switching costs and regulatory interventions that shift the features of the minimum standard.

[29]Using the Medicare conventions, each plan here is defined as a plan-region combination, or to be more precise as a contract id and plan id combinations. This implies that a plan of the same insurer with the same name, but e.g. different prices in different regions will be counted as a different plan here.

1.5.2 Quantifying the effect of switching costs on the allocation of risks

The estimated model of contract choice in Medicare Part D allows me to explore how switching costs are altering the allocation of risks among contracts over time. This counterfactual exercise is similar in spirit to the analysis in Handel (2013) in the employer-provided insurance setting. I begin with a discussion of two stylized examples. These examples illustrate, first, that whether switching costs ameliorate or exacerbate adverse selection depends crucially on the evolution of the relative generosity and the relative prices of contracts in relation to the contract characteristics at the first choice incidence. Second, these examples suggest that accounting for changes in the contract space beyond price adjustments may have important implications for the conclusions we draw about the interaction between adverse selection and switching costs. In particular, in insurance settings akin to Medicare Part D, insurers can control when they let contracts enter the market and can change their contracts on a variety of dimensions in addition to premiums. These margins of adjustment[30] are crucial for the distribution of risks and the interaction with switching costs. Moreover, they imply that in an environment where adjustments to the contract space are undertaken simultaneously by a large number of insurers along several different dimensions, the net effect of the switching costs on the allocation of risks is an inherently empirical question. Indeed, the simulation results for the evolution of the full contract space in Medicare Part D presented after the stylized examples suggest that, on net, costly switching helped support an adversely selected equilibrium in this environment.

Stylized examples of switching costs altering risk sorting within a simplified contract menu I provide two stylized examples using a significantly simplified version of the Part D contract space to illustrate the opposite effects that switching frictions may have on selection. Each example considers a menu of three contracts of different types and simulates the evolution of choices and risk selection among these three contracts, assuming that they were the only choices available to all individuals on the market for four years. The two different menus I consider were actually observed on the market. For each, I do the simulation exercise twice - assuming either costly or costless switching. For the simulations, I use the estimated parameters of the model, including the switching cost parameter, the observed demographics, and the risk scores for all individuals in the estimation sample.

Consider the contract menu in the first example. In this menu, the insurer offers *Type 1* and *Type 2* plans in the first year of the market, adding a *Type 3* plan in the second year. Panel A of Figure 1-9 plots the development of the relative premium between the *Type 3* and *Type 1* plan for the three years when both are available. The relative premium is quite low in the first year *Type 3* contract is introduced; nevertheless, in the scenario with costly switching very few individuals take up this more comprehensive plan. This is in stark contrast to the scenario with costless switching, in which the enrollment in the *Type 3* plan jumps to 36%. Importantly, we see that with the jump in the enrollment share, costless switching also results in adversely selected enrollment in the *Type*

[30]The empirical model in the current paper takes the decisions of insurers along these margins as given, except for the cases where I impose counterfactual adjustments in simulating minimum standard policies. Endogenizing the entry and contract choice decisions is a highly interesting direction for future research.

3 plan relative to *Type 1*. With costless switching, selection gets more acute in 2009, as the relative premium of the contract increases and enrollment drops significantly. In the scenario with switching costs, neither the enrollment, nor the relative risk change much over time. The enrollment in *Type 3* contract stays at below 5% level and the relative risk stays substantially below the relative risk in the scenario with costless switching. This setting thus illustrates the natural intuition that switching costs should be muting adverse selection.

Now consider a different contract menu example. In this example, the insurer offers a generous *Type 4* plan priced at a substantial premium over the *Type 1* plan in the first year, which has a very low premium. *Type 4* plan collects the highest risks on the market in year one. In year two, the insurer demotes *Type 4* plan to be *Type 3* and raises its premium, while also raising the premium of *Type 1* plan. Figure 1-9 illustrates that the relative premium between the most and the least generous plans of this insurer first rises and then falls over time. The same figure illustrates the simulated enrollment and risk-sorting patterns for this contract menu. We see that in this setting, where a very generous plan collected the highest risk in the first year, switching costs support the large differential in risk between the most and least comprehensive plan over time. The simulation of choices suggests that the difference in risks would have been lower in a counterfactual with costless switching. In other words, in this scenario of contract menu evolution, switching costs exacerbate, rather than mute, selection.

Simulation of risk allocation with costly and costless switching on the full contract space Table 1.10 illustrates the results of the counterfactual analysis that completely shuts down the switching cost channel in the choice model. I allow for two scenarios in the simulation. One scenario takes the observed contract premiums as given; the other scenario allows insurers to endogenously adjust premiums to the new sorting patterns according to the contract pricing model described in Section 1.4.2. Both scenarios take the other features of the contracts, as well as the observed individual demographics and expected spending, as given. Table 1.10 focuses on four moments of the data aggregated at three-type plan level - enrollment shares, average expected spending, average ex-post spending, and average premiums. All results are aggregated to the three types of plans. I present the results for year 2009, which serves as a summary statistic of the market development with switching costs since 2006.

I perform three simulations of the model. The first one, marked with A, simulates the model with estimated switching costs, taking the observed prices and pre-2009 choices as given.[31] This step creates a baseline that takes into account the simulation error and which I use instead of the actual observed 2009 outcomes as a comparison benchmark in analyzing the scenarios without switching

[31] As part of specification checks, I tested the alternative approach of simulating the model from 2006 onward rather than taking the observed lagged choices as given in 2009. While the simulation error accumulates starker over several simulation periods in this case, this doesn't change the analysis in a substantive way. The baseline approach pursued in the main text renders itself better to the interpretation of the switching cost reduction as a sudden policy shock in one year. This point is irrelevant for the simulations without switching costs, since lagged choices do not enter the utility function.

costs.[32] Baseline enrollment shares, risk sorting, and prices paid, closely resemble the descriptive evidence presented earlier. *Type 2* plans with reduced deductible and no gap coverage have the highest enrollment share - 72% and about average (in-sample) risk profile of $1,926 expected annual spending. *Type 3* plans with partial gap coverage are adversely selected with the average expected spending of $2,368. The difference in the average risk between the *Type 1* and *Type 3* plans is substantial - expected spending in type 3 plans is 22% or $526 higher. *Type 3* plans have annual premiums that are on average $400 higher than in plans without partial gap coverage. Enrollment in *Type 3* plans is relatively low - at 9%.

The next simulation, marked with B, sets the switching cost parameter γ in the utility function to zero for all individuals, which eliminates the inertia channel. This simulation shows what enrollment shares and risk sorting would look like absent choice frictions if insurers kept their prices at the observed 2009 levels.[33] Two key changes are notable relative to the baseline with switching costs. First, there is some change in enrollment shares, as with costless switching individuals can actively respond to adjustments in the price and other characteristics of contracts relative to 2006. Enrollment in *Type 1* plans increases by 5 percentage points, as enrollees respond to lower premiums available for these plans. In this simulation, enrollees pay on average $57 a year less for their *Type 1* plan. Respectively, the enrollment in *Type 2* and *Type 3* plans decreases. Beneficiaries also select cheaper contracts within these types of plans, paying on average $27 and $75 less in annual premiums. Second, the simulation predicts smoother distribution of the average expected spending among plans. Both *Type 1* and *Type 3* plans experience changes in average risks that move them closer to the average expected spending in the sample. This decreases the relative average risk between these plans from $526 at baseline with switching cost to $463 in the counterfactual without switching costs and fixed prices - a decrease of 12%. Thus, even before allowing the insurers to respond to the change in sorting patterns by adjusting their prices (which we would expect to amplify these effects), we see that removing switching costs mutes adverse selection in this setting in the sense of balancing the distribution of risks across plans.

In the last simulation, I allow insurers to respond to the changes in the sorting patterns induced by the removal of switching costs. Insurers are assume d to adjust their premiums according to the stylized pricing model outlined in Section 1.4.2 in those years where the absence of switching costs changes sorting across plans.[34] Allowing prices to adjust to the sorting without switching costs amplifies the results in B. The most generous *Type 3* plans with partial gap coverage become $100

[32] As discussed in Section 1.5.1, the simulated baseline is close to the observed moments in the data.

[33] While this scenario may appear very unrealistic for a competitive market and its primary purpose is to illustrate the separate effects of risk-sorting response and supply-side adjustments, the regulatory environment in Medicare Part D is such that no changes in the prices of contracts from the individual's point of view is in fact conceivable. Since CMS regulates how bids by insurers get translated into beneficiary premiums, adjusting this mechanism so as to freeze relative prices faced by individual beneficiaries in addition to policies that reduce switching costs is, in principle, possible.

[34] This implies that I keep observed prices in 2006 and 2007. Since premiums are based on lagged sorting, in these years switching cost counterfactuals should not be affecting prices. I simulate prices for years 2008 and 2009, using the simulated changes in beneficiary choices due to costless switching in 2007 and 2008 respectively.

cheaper in response to lower expected risk, which triggers an increase in their enrollment share relative to the simulation in B. Risk sorting among plans becomes even less acute and all of them move even closer to the average. The relative average risk between the *Type 1* minimum standard plans and *Type 3* plans with partial gap coverage falls to \$414, which is a 21% decrease in the risk difference relative to the baseline simulation. Figure 1-10 plots the counterfactual risk CDFs by the type of plan for this simulation. The graph demonstrates that the decrease in risk difference holds throughout the whole distribution of risks and not only for the mean.

1.5.3 Effect of minimum standard regulation on the allocation of risks with costly and costless switching

The analysis in the preceding section is consistent with the idea that switching costs significantly alter the response of risk-sorting to changes in insurance contracts. It also demonstrates that the effect of switching costs on adverse selection depends critically on the exact evolution of the contract space relative to the initial conditions. For public health insurance environments, this insight has important policy implications. Multiple regulatory interventions that directly affect the contract space are ubiquitous in such environments. The fact that switching costs alter the response of risk-sorting to contract changes implies that switching costs will also significantly alter how the regulatory interventions impact the allocation of risk. In the following simulations, recorded in Table 1.11, I explore this hypothesis on the example of the Part D's minimum standard regulation. I first quantify how changes in the minimum standard policy impact risk sorting across contracts in the status quo with costly switching. I then consider how switching costs are altering this impact. I find two surprising patterns. First, I find that some minimum standard reforms may have an unintuitive effect on the level of selection and insurance coverage purchased by individuals. Specifically, my simulations suggest that lowering the minimum standard would, in fact, lead beneficiaries to purchase more comprehensive coverage. Second, I find that switching costs completely mute the ability of the minimum standard regulation to change the distribution of risks across contracts.

The first counterfactual I consider, quantifies the effect of a local deviation in the minimum standard level. In this counterfactual, I change the 2009 minimum standard to the 2006 level. This implies a reduction in the deductible offered by all *Type 1* plans in 2009 from the observed \$295 to \$250, and a reduction in the ICL from the observed \$2,700 to \$2,250 for all three types of plans. This exercise simulates the short-run effects of a counterfactual policy shock in 2009. Ex ante, we would expect that absent choice frictions and price adjustment, *Type 1* enrollment share should increases and its risk pool worsen, since *Type 1* contracts with a lower deductible at the same prices become more attractive relative to *Type 2* contracts. With endogenous re-pricing of contracts, the counterfactual policy has a theoretically ambiguous effect. To test which direction the empirical simulation takes, I consider three scenarios of the choice environment in this policy experiment. One with switching costs and observed prices, another without switching costs, but still observed prices kept fixed, and lastly the no switching cost scenario where I allow insurers to adjust prices to the new sorting and regulatory conditions. Note that I do not allow for any other

changes in the contract space, and take the entry and exit of plans, as well as the changes of other characteristics between 2008 and 2009 as given. The effect of this policy deviation is small, but follows the predicted direction. In the baseline scenario that keeps switching costs and observed prices, *Type 1* enrollment share moves from 18% to 19%. The effect on enrollment is only somewhat starker in the other two scenarios where I shut down the inertia channel. The local deviation in the minimum standard policy considered also has no economically meaningful effect on risk sorting across different types of plans even with frictionless switching.

The second simulation implements a counterfactual policy of loosening the minimum standard. Specifically, I analyze the case where the standard defined benefit would have been a high-deductible plan, which is a policy that had been considered by the government as a possible option during the regulatory design of Medicare Part D. I simulate the market outcomes for a policy shock that forces all *Type 1* plans to have $\$1,000$ deductible in 2009, but keeps other features of *Type 1* and other contract types the same. I again consider three scenarios: with switching costs at observed prices and without switching costs with and without price adjustment. Not surprisingly, the response to this policy is much starker than to the previous local deviation in the minimum standard. In the high deductible scenario, *Type 1* plans loose a substantial share of their enrollment. It drops from 19% to 10% in the scenario with switching costs and from 23% and 26% to 12% and 20% in the scenarios with no switching costs with and without price adjustment respectively. The effect on the allocation of risk is small, even in scenarios that completely shut down the inertia channel and re-price the contract. This is not too surprising, given that in the reduced-form evidence of adverse selection we found relatively little selection on the deductible margin. Although the observation that *Type 1* plans substantially loose their enrollment share under observed prices is natural, since increasing the deductible makes the plans much less generous at the same price, the full equilibrium result is less intuitive. We could have expected that with a price adjustment that makes the high deductible plan significantly cheaper than the next available alternative (the simulated annual premium drops to $45), there would exist substantial demand for such a "catastrophic coverage" plan. Evidently, for the beneficiaries whose observed choices were used to estimate the model and who selected into Part D insurance to begin with, the valuation of the first-dollar coverage is very high. In this pool of beneficiaries, introducing a high deductible plan amplifies the tendency to a pooling equilibrium on the contracts with low deductible - a tendency that has been observed in this program over time as the regulator has been increasing the SDB deductible every year. Although not modeled in the current setting, it is possible that the existence of a relatively cheap high deductible plan would attract more individuals on the extensive margin of the program, potentially improving the overall welfare in Part D.

The last policy experiment that I consider corresponds to the changes in the Part D minimum standard regulation envisioned under the Affordable Care Act (ACA). One provision of the ACA is that the standard plan in Part D will no longer have the coverage gap. Beyond any political economy reasons, an economic rational for the introduction of this change in the minimum standard is the unraveling of plans with full coverage in the gap on the private market. In this counterfactual

simulation I force all plans to have the level of 2009 catastrophic coverage threshold ($6,145) as their initial coverage limit (originally $2,700). This implies that *Type 2* and *Type 3* plans become identical in terms of their key financial characteristics, as both of them have reduced deductible and are imposed to have the same coverage limit. The counterfactual simulations suggest that switching costs play a key role in determining the effect of this policy change on risk allocation. Ex ante, we would expect that, since there is no gap coverage dimension of differentiation among plans anymore, this policy should eliminate the acute selection that took place on the gap coverage margin. Indeed, my results suggest that exactly this effect takes place, but only in the scenarios, where I shut down the inertia channel. In these scenarios, all types of plans converge to having practically identical pool of risks of around $1,950 in expected spending. In the scenario with costly switching, however, the simulated policy does not change the distribution of risks. Figure 1-11 plots the counterfactual risk CDFs for this simulation, demonstrating that the result holds throughout the whole distribution of risks and not only for the differences in mean risk.

1.6 Conclusion

In this paper I have documented evidence of adverse selection and switching costs in a highly regulated Medicare Part D environment using a parsimonious classification of the contract space and detailed administrative data. I have also shown that switching costs have important interactions with the adverse selection in this complex health insurance program. In particular, I have shown that in this environment the initial conditions led switching costs to support an adversely selected equilibrium over time, in the sense that different types of plans would have had more similar average risks if switching were costless. In considering an important channel that drives the changes in Part D contracts over time - the minimum standard policy, I found that the empirical effects of changing such regulation may be unexpected. For instance, my simulations suggest that loosening the minimum standard could, in fact, increase the amount of coverage that individuals purchase. Further, I have shown that the market imperfection of costly switching plays an important role in determining the risk-sorting outcomes of this policy intervention, which is targeted at correcting a different market failure - adverse selection. In particular, I find that in the presence of switching costs, tightening the minimum standard requirement by "filling" the donut hole is unlikely to have the intuitive effect of balancing risks across different contracts.

More broadly, this paper argues that in considering the polices that may improve consumer choice by eliminating or decreasing switching frictions in the increasingly common public health insurance settings with regulated competition, we have to take into account the nuanced interconnections of the different market imperfections with the regulatory instruments targeted at correcting them. The caution, of course, comes from the caveat that in this work I represented the reaction of insurers to policy changes in a very stylized way. Expanding the model to allow insurers to endogenously react to adverse selection and regulation in their choices of contract characteristics in a competitive setting provides a fruitful area for future research.

Figure 1-1: Minimum standard policy in Medicare Part D: Standard Defined Benefit in 2006-2009

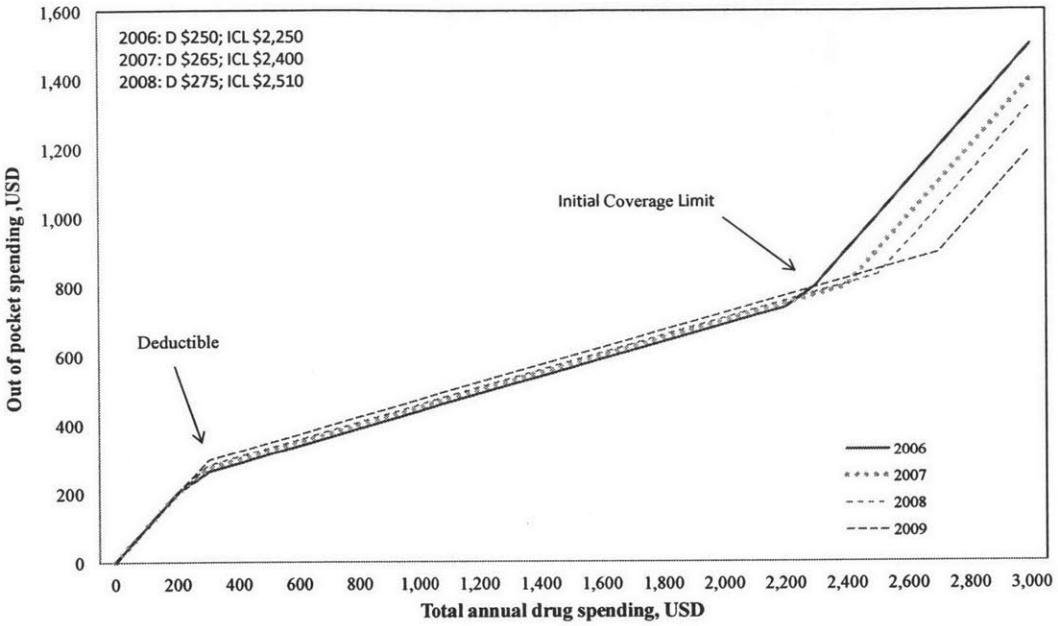

Standard Defined Benefit parameters as set by Medicare in years 2006-2009:

	2006	2007	2008	2009
Deductible	$ 250	$ 265	$ 275	$ 295
Initial Coverage Limit (ICL)	$ 2,250	$ 2,400	$ 2,510	$ 2,700
Maximum OOP before catastrophic coverage starts	$ 3,600	$ 3,850	$ 4,050	$ 4,350

Insurers in the Medicare Part D program are required to provide coverage that gives at least the same actuarial value as the Standard Defined Benefit (SDB). The SDB design features a deductible, a co-insurance rate of 25% up to the initial coverage limit (ICL) and the subsequent "donut hole" that has a 100% co-insurance until the individual reaches the catastrophic coverage arm of the contract. The graph illustrates these features of the SDB by mapping the total annual drug spending into the out-of-pocket expenditure. As the figure illustrates, the generosity of the SDB changed over time.

Figure 1-2: Evolution of premiums, claims and enrollment shares by type of plan: Types 1 and 2

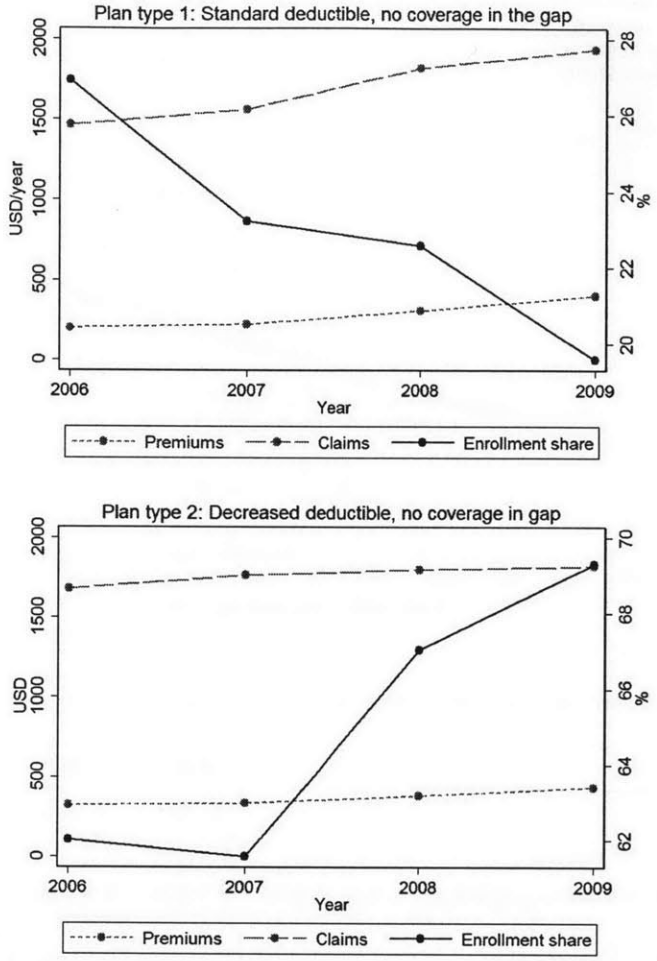

The data in the panels is based on the baseline sample for years 2006-2009. The calculation of average premiums and claims is enrollment-weighted. Claims include the total annual drug spending, without accounting for cost-sharing. To ensure comparability of the observations, the spending of individuals with enrollment shorter than 12 months (primarily 65 year olds) was extrapolated to the full year. The panel for *Type 4* plan illustrates the significant drop in enrollment and contemporaneous increase in the premiums and claims in the plans with full coverage in the gap. The claims and premiums for this plan type are recorded on log-scale for visual clarity.

Figure 1-3: Evolution of premiums, claims and enrollment shares by type of plan: Type 3 and 4

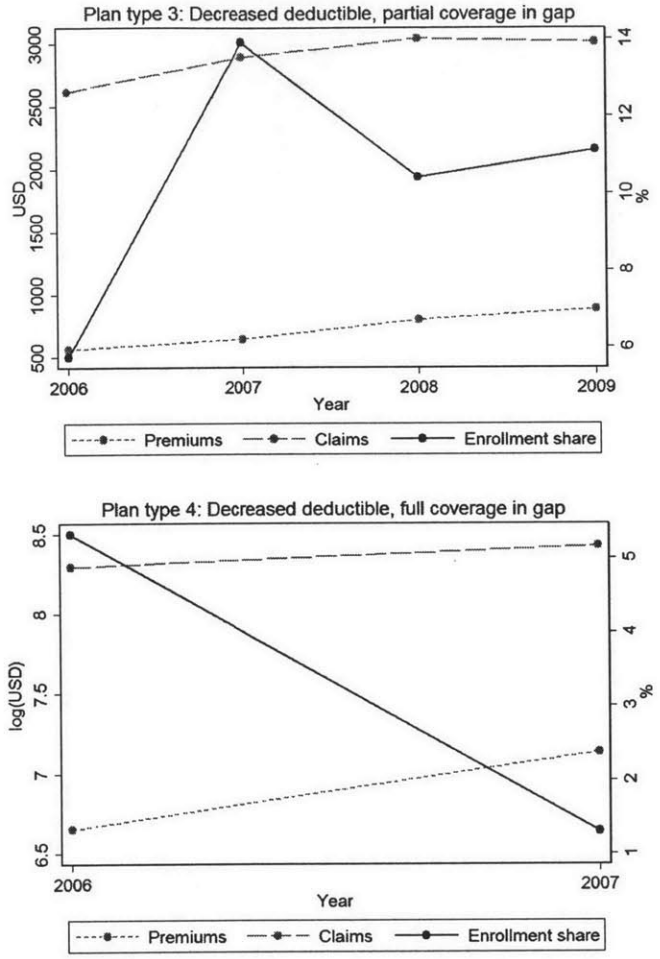

The data in the panels is based on the baseline sample for years 2006-2009. The calculation of average premiums and claims is enrollment-weighted. Claims include the total annual drug spending, without accounting for cost-sharing. To ensure comparability of the observations, the spending of individuals with enrollment shorter than 12 months (primarily 65 year olds) was extrapolated to the full year. The panel for *Type 4* plan illustrates the significant drop in enrollment and contemporaneous increase in the premiums and claims in the plans with full coverage in the gap. The claims and premiums for this plan type are recorded on log-scale for visual clarity.

Figure 1-4: Positive correlation tests for the presence of asymmetric information

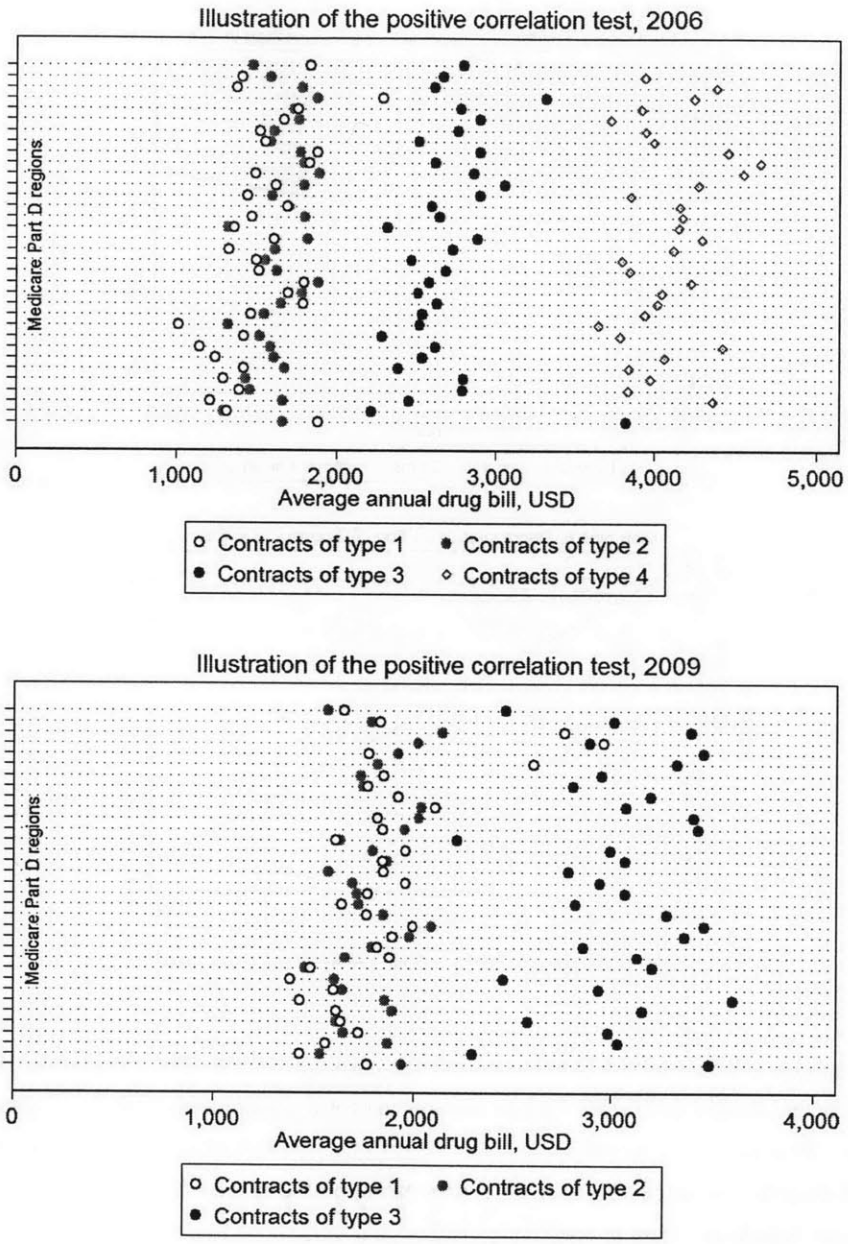

Each dot plots the average annual total drug expenditure observed in years 2006 and 2009 for enrollees in different types of Medicare Part D plans using the baseline sample. The averages are calculated separately within each region.

Figure 1-5: Evidence of adverse selection: distribution of risks by type of plan using risk score measure

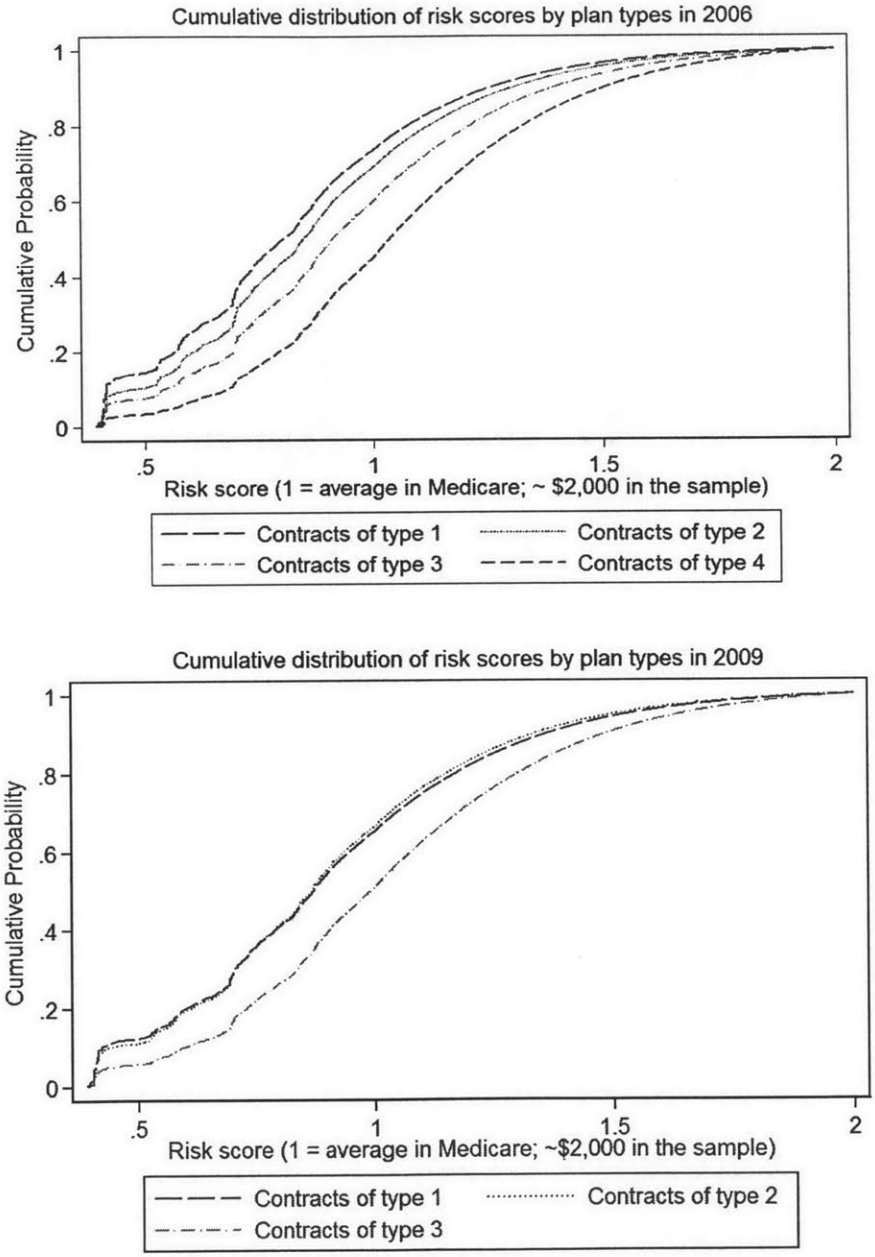

The graphs plot empirical CDFs of ex-ante risk in plans of different generosity types in years 2006 and 2009. The risk is measured using individual risk scores that are based on lagged diagnostic information and thus do not suffer from the moral hazard concern in drug expenditures.

Figure 1-6: Evidence of switching costs: over time development of insurers' enrollment shares for new and continuing beneficiaries

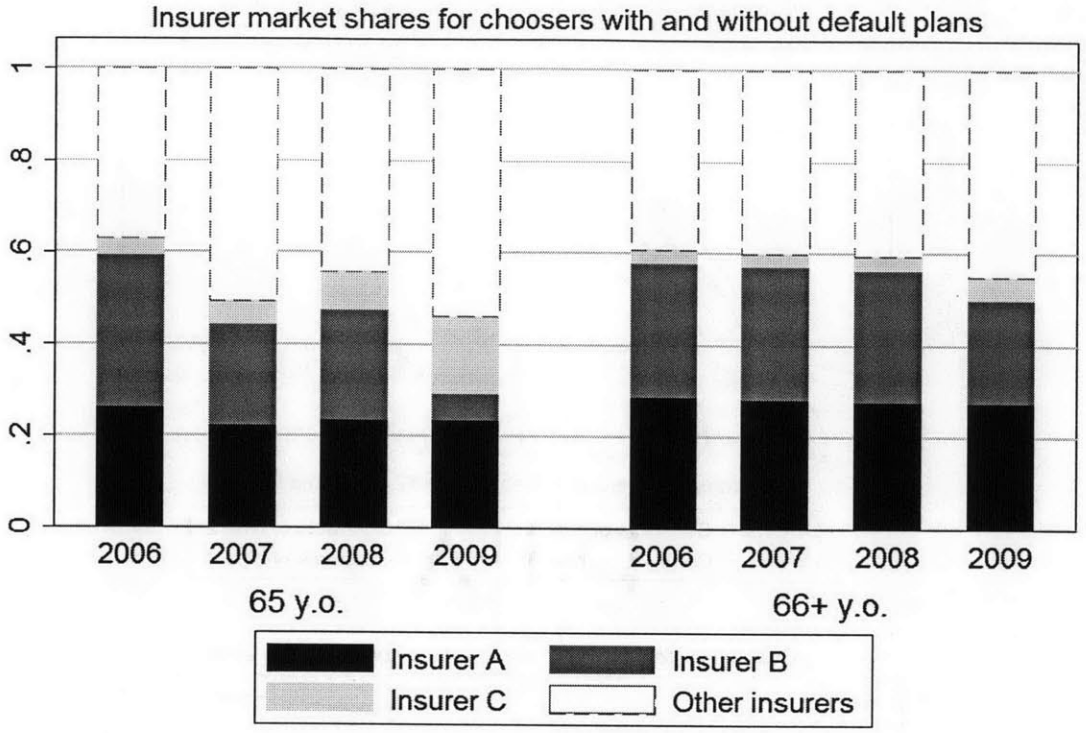

Based on the working sample, may not coincide with agggregate market reports

The graph uses data of the baseline sample. We observe that the choices of the 65 year old individuals newly entering the program, who by definition do not have incumbent plans, are much more volatile and responsive to the market conditions over time, than choices of the individuals in the existing cohorts, who usually have the default option of their incumbent plan available. Insurer "identities" here are constructed using contract encryption in the administrative data. Because of the data encryption, separate insurers may have been identified with error. The corresponding real names of the insurance companies are not known to the researcher.

Figure 1-7: In-sample fit of the choice model

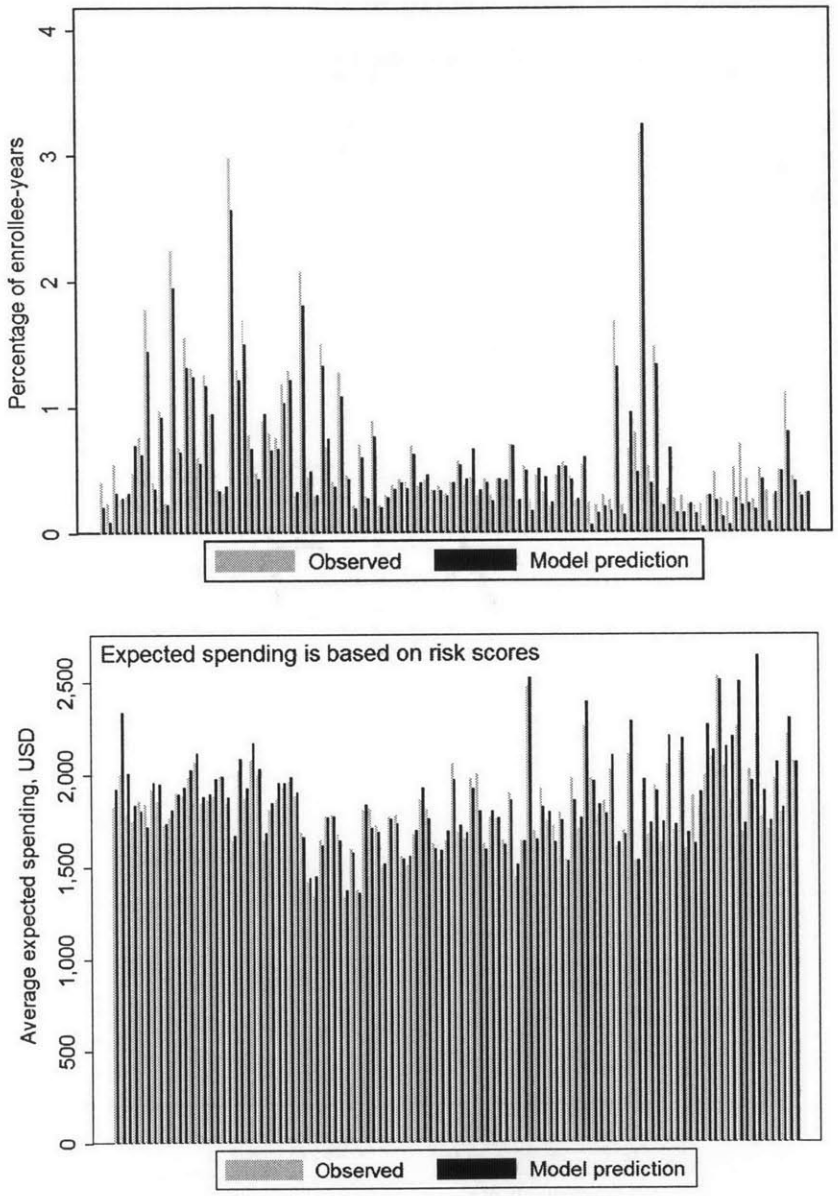

To construct simulated enrollment, the estimated coefficients of the choice model together with simulated random components were used to find the contract with the highest utility in each individual's choice set. The observed risk scores of the individuals predicted to enroll in different plans were used to compute the average predicted risk. Each pair of bars in the graph represents a different Medicare Part D plan ("plan" is region-specific). The graphs display only top 90 out of 2,357 contracts.

Figure 1-8: Key descriptive evidence generated in the model: empirical distribution of expected spending by type of plan over time; in-sample data

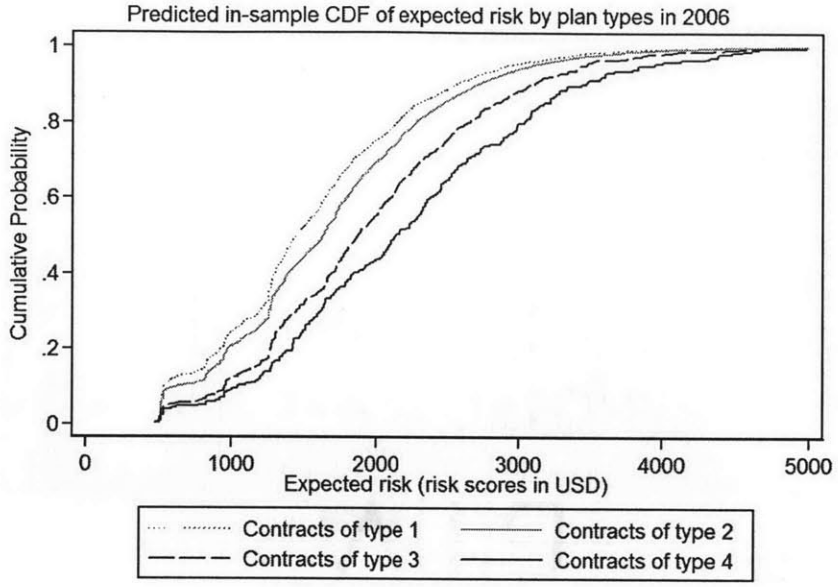

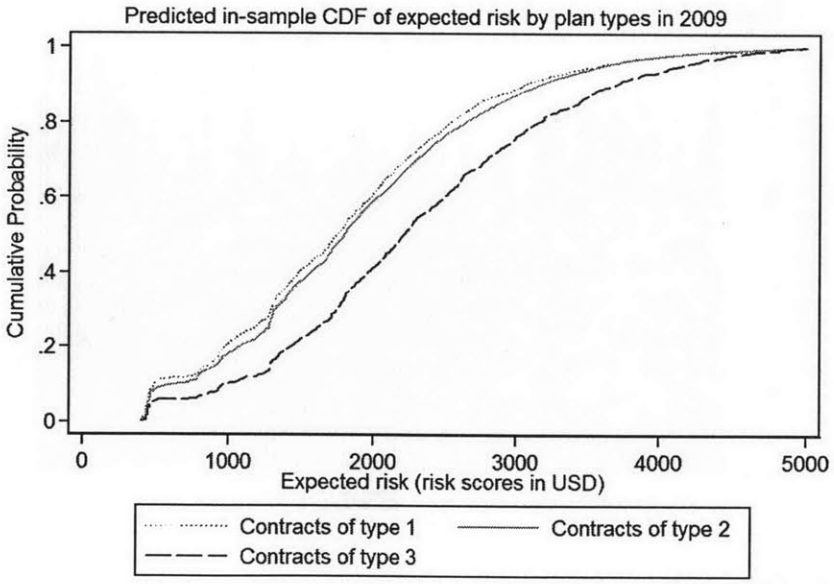

Figure 1-9: Stylized examples of interaction between switching costs and adverse selection

- *Panel A: Example of switching costs muting adverse selection*

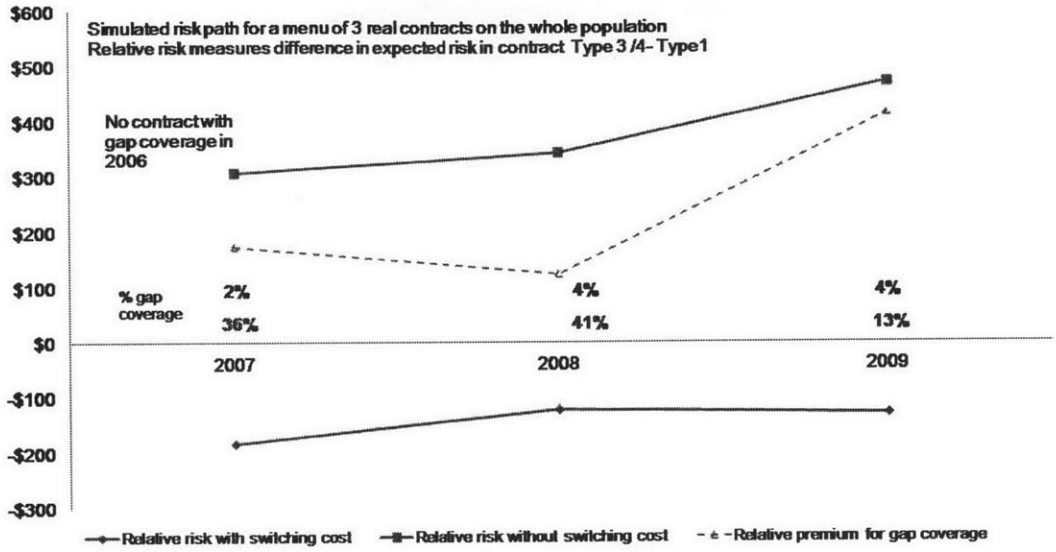

- *Panel B: Example of switching costs supporting adverse selection*

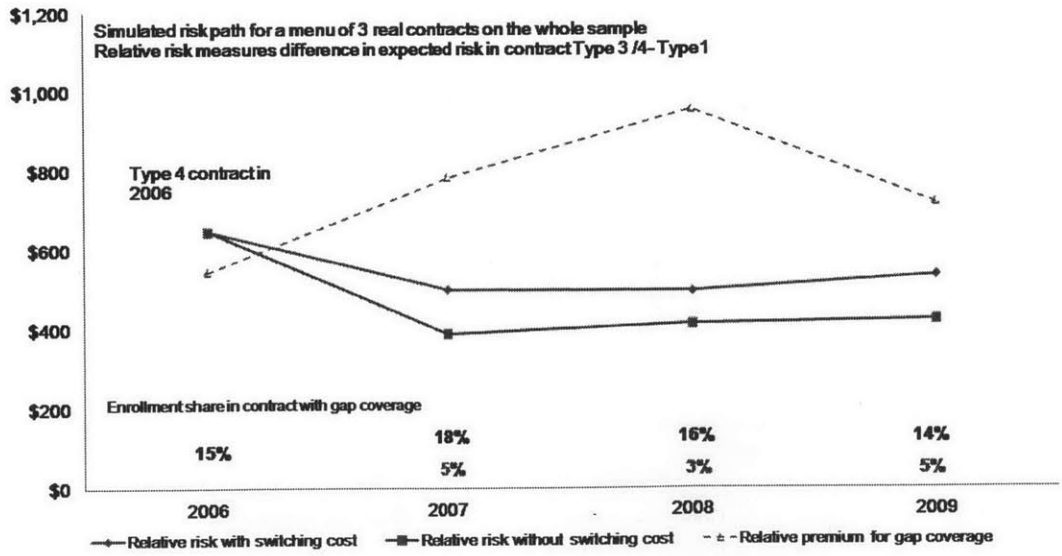

The graphs illustrate the relative risks and enrollment choices simulated using the estimated parameters of the model with and without switching costs in two examples of stylized Part D contract space. Section 1.5.2 discusses how these examples were constructed and their interpretation.

Figure 1-10: Counterfactual risk allocation without switching costs: baseline versus counterfactual distribution of risks by contract type in 2009

- *Panel A: Baseline simulated distribution of risks*

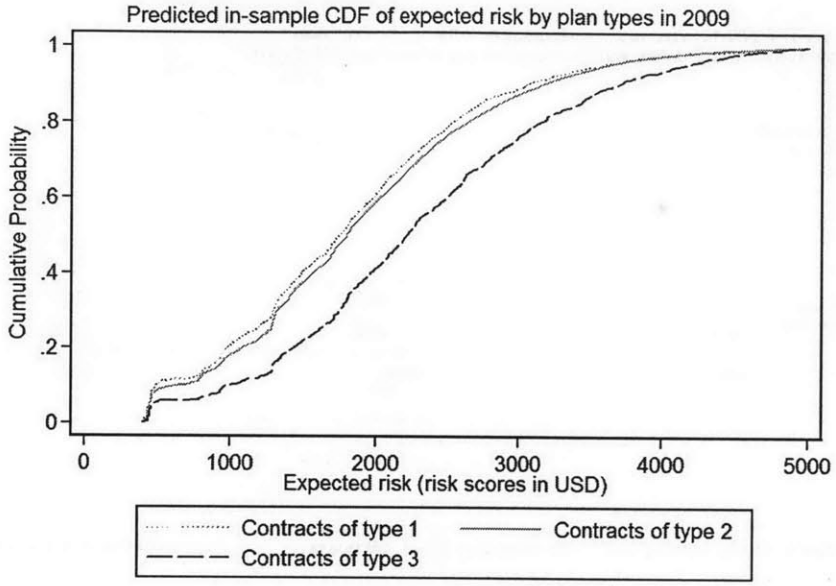

- *Panel B: Simulated distribution of risks without switching costs and endogenous re-pricing*

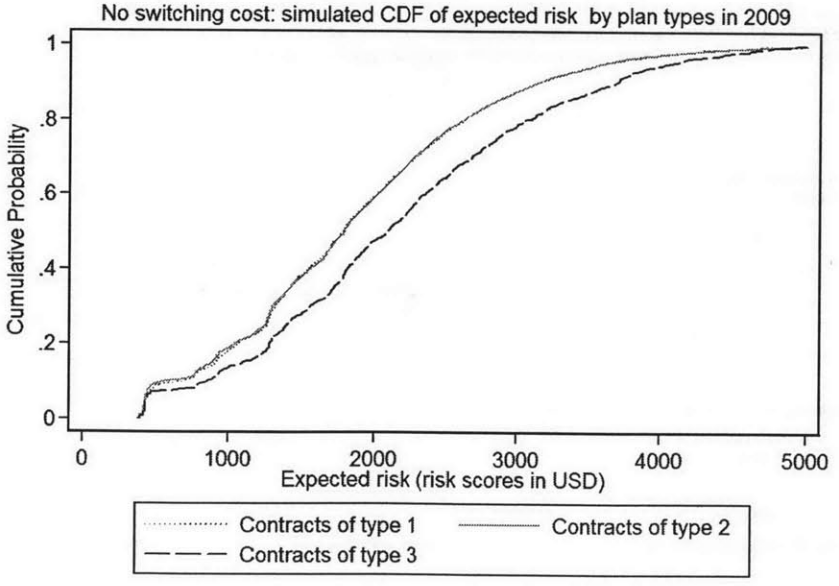

Figure 1-11: Minimum standard counterfactual: the role of switching costs in determining the effect of the policy on the distribution of risks among contracts

- *Panel A: Distribution of risks* ***with*** *switching costs; "filling the donut hole" counterfactual*

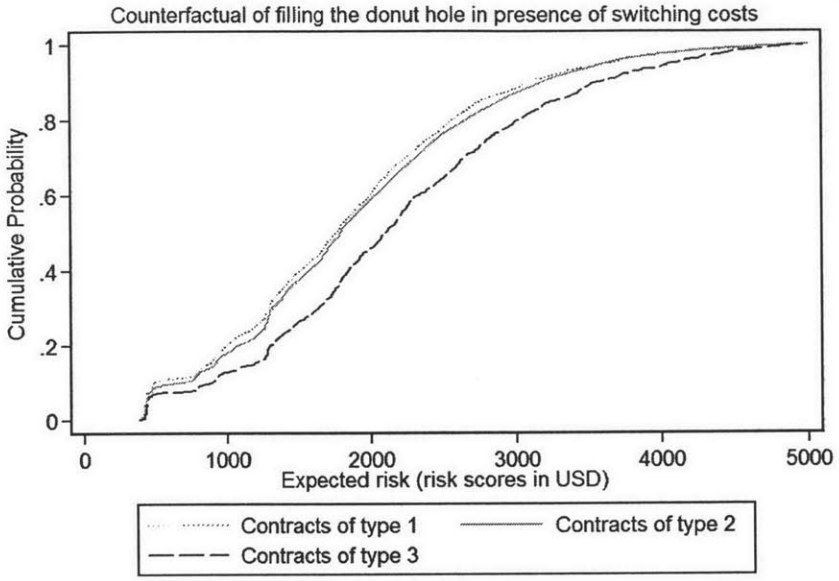

- *Panel B: Distribution of risks* ***without*** *switching costs; "filling the donut hole" counterfactual*

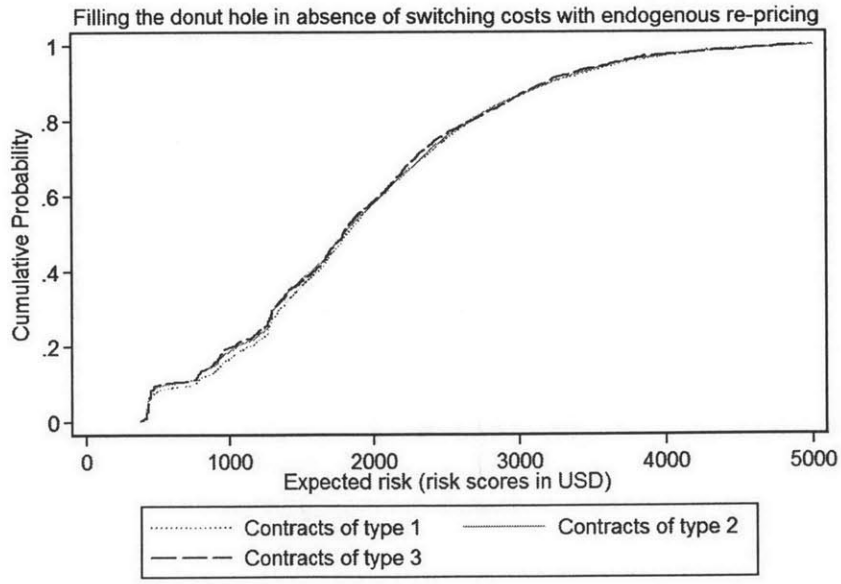

Table 1.1: Summary statistics of the full sample, baseline sample and panel sub-sample

	Full sample	Baseline sample	Panel sub-sample
2006			
N	9,999,538	1,221,252	871,818
Age (std.dev)	72 (12)	76 (8)	75 (7)
Female	0.59	0.65	0.65
White	0.84	0.95	0.96
ESRD	0.01	0.003	0.001
Risk score (σ)	n/a	0.89 (0.34)	0.86 (0.31)
Annual drug spending (σ)	n/a	1,518 (1,899)	1,449 (1,733)
2007			
N	10,176,611	1,307,966	911,403
Age (σ)	72 (12)	76 (8)	75 (7)
Female	0.59	0.63	0.65
White	0.84	0.95	0.96
ESRD	0.01	0.003	0.002
Risk score (σ)	n/a	0.90 (0.35)	0.88 (0.32)
Annual drug spending (σ)	n/a	1,883 (2,407)	1,832 (2,227)
2008			
N	10,369,814	1,356,861	954,494
Age (σ)	72 (12)	76 (8)	76 (7)
Female	0.58	0.63	0.65
White	0.83	0.95	0.96
ESRD	0.01	0.003	0.002
Risk score (σ)	n/a	0.91 (0.36)	0.90 (0.34)
Annual drug spending (σ)	n/a	1,907 (2,648)	1,869 (2,479)
2009			
N	9,781,213	1,365,239	998,014
Age (σ)	71 (12)	76 (8)	76 (8)
Female	0.55	0.63	0.65
White	0.83	0.95	0.96
ESRD	0.01	0.003	0.003
Risk score (σ)	n/a	0.92 (0.36)	0.92 (0.35)
Annual drug spending (σ)	n/a	1,950 (2,973)	1,947 (2,934)

Table 1.2: Evidence of adverse selection: positive correlation tests using realized ex-post drug expenditures and diagnosis-based free risk scores

$$Y_{irt} = \alpha_r + \delta_t + \sum_{k=2}^{k=4} \beta_k \mathbf{1}\{ContractType_{it} = k\} + \epsilon_{irt}$$

	(1) Annual drug spending	(2) Risk score	(3) Risk score projected to USD
Contracts of type 1	reference category		
Contracts of type 2	-2.047	0.00927	24.98
	(71.83)	(0.00730)	(20.89)
Contracts of type 3	1213.4***	0.146***	415.0***
	(105.5)	(0.0107)	(30.46)
Contracts of type 4	3081.3***	0.260***	728.5***
	(70.71)	(0.00521)	(14.57)
Year FE	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
N	3,892,280	3,892,280	3,892,280
Mean Y	1948.7	0.920	1948.7
St. dev. Y	2712.2	0.357	1018.7

Standard errors in parentheses clustered at the region level

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

The regressions test for the correlation between the generosity of the chosen insurance contract and the measures of individual risk conditional on geographic regions, which are used for the pricing of contracts. The results are based on the pooled data of the baseline sample that spans years 2007-2009. Details of contract classification typology are discussed in the main text. Outcome in column (1) is measured as annual recorded ex-post drug spending without accounting for cost-sharing. Outcome in column (2) are risk scores based on diagnostic information from Medicare A/B (1=average risk in Medicare). In column (3) the outcome variable is the predicted value from the regression: $E[AnnualDrugSpending_i|.] = \alpha + \beta RiskScore_i$ run separately for each cross-section. The specification in column (3) thus considers the part of drug spending variation that is attributable to the observed medical diagnosis and thus forecastable risk and doesn't contain moral hazard effects. The results in column (3) can therefore be interpreted as the lower bound of the adverse selection portion of the overall level of asymmetric information.

Table 1.3: Evidence of switching costs: share of enrollees that choose their "default" plan in a given year

	2007	2008	2009
A. All plans			
Probability of choosing default plan for 66+ y.o. enrollees	89.9 %	88.7 %	89.1 %
N	1,089,978	1,162,545	1,194,036
B. Default plan in t was re-classified in its type vs $t-1$			
Probability of choosing default plan for 66+ y.o. enrollees	73.4 %	78.9 %	87.9 %
% of sample	14 %	7 %	4 %
C. Default plan in t had the same type as in $t-1$			
Probability of choosing default plan for 66+ y.o. enrollees	92.6 %	89.5 %	89.3 %
% of sample	86 %	93 %	96 %

The table shows the share of individuals in each year that choose to remain in their "default" plan among those for whom default plans could be defined. The data is a sub-sample with two-year panel observations from the baseline sample of enrollees. The observations do not include 65 year olds, as these by definition do not have default plans. A default plan flag is constructed by using plan cross-walks provided by CMS and requires observing the policy in which the individual was enrolled in two consecutive years. The default plan enrollment flag is set to equal one in two cases. First, and most common, if the individual enrolled in the plan with exactly the same CMS id in year $t-1$ and t. Second, if in year t individual enrolled in a plan that is not identical to the plan id in $t-1$, but was recorded as a plan consolidating the original $t-1$ plan, in which case the CMS policy is to default the individual into the consolidating plan if the individual takes no action of choosing a different option. Panels B and C define plan "type" according to the 4-type topology used in the main text. These panels divide the sample of beneficiaries into those who did and did not experience a significant supply-induced change in their incumbent plan. Panel B excludes observations for individuals whose plans were terminated and no default option was available. These individuals comprise less than 1% of the sample.

Table 1.4: Evidence of switching costs: choice patterns in 2006-2009 aggregated to 4 types of plans tracked for cohorts entering in different years

Cohorts of 65 year olds whose chosen plans were not re-classified into a different type by the insurer throughout all tracked years										
	65 y.o. in 2006				65 y.o. in 2007			65 y.o. in 2008		65 y.o. in 2009
A. *Enrollment shares*	2006	2007	2008	2009	2007	2008	2009	2008	2009	2009
Contracts of type 1	22 %	22 %	19 %	17 %	17 %	15 %	14 %	10 %	11 %	12 %
Contracts of type 2	73 %	73 %	77 %	79 %	72 %	75 %	77 %	82 %	82 %	82 %
Contracts of type 3	4 %	5 %	5 %	4 %	11 %	11 %	10 %	7 %	7 %	5 %
N	37,500	37,500	37,500	37,500	35,759	35,759	35,759	40,960	40,960	43, 520
B. *Incremental premium*	**in year 2006**				**in year 2007**			**in year 2008**		**in year 2009**
Contracts of type 2	\$138				\$125			\$54		\$37
Contracts of type 3	\$375				\$360			\$410		\$469

Panel A shows enrollment shares in each year across three types of plans for cohorts of 65 year olds entering the program in different years. The sample includes only individuals, whose default plans were not re-classified in their type by the supply-side throughout the observed enrollment time. The choices of a given cohort are recorded for all subsequent years of available data. The calculation is based on the panel sub-sample data, as described in the data construction appendix. The table shows raw enrollment shares as observed in each year subject to the classification of contracts into the 4-type topology. The choices of cohorts show persistence over time and differ from the choices of newly entering cohorts within the same year. The difference is especially apparent between the first two and last two years of the data. Panel B adds information about the development of the relative annual premiums over-time. The premiums are reported as increments relative to the *Type 1* contract with SDB deductible and no coverage in the gap. To reflect the market conditions, the premiums are constructed as averages weighted by enrollment of all 65 year olds in the respective years.

Table 1.5: Evidence of switching costs: price sensitivity estimates for individuals with and without default plans

	Age of beneficiaries					
Price coeff. [p-value]	65	66	67	68	69	70
	Baseline	Interaction	Interaction	Interaction	Interaction	Interaction
2006	**-0.003**	0.0001	0.0002	0.0006	-0.0001	0.0006
	[0.000]	[0.809]	[0.683]	[0.386]	[0.876]	[0.321]
2007	**-0.003**	**0.0018**	**0.0012**	**0.0011**	**0.0013**	**0.0010**
	[0.000]	[0.002]	[0.035]	[0.031]	[0.002]	[0.040]
2008	**-0.003**	**0.0022**	**0.0023**	**0.0021**	**0.0019**	**0.0020**
	[0.000]	[0.000]	[0.000]	[0.000]	[0.001]	[0.001]
2009	**-0.010**	**0.0072**	**0.0085**	**0.0090**	**0.0085**	**0.0084**
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]

The price coefficients are estimated using the following random utility specification: $u_{ij} = -\alpha_{65}p_{ij} + \alpha_{66}p_{ij}\mathbf{1}\{Age = 66\} + \alpha_{67}p_{ij}\mathbf{1}\{Age = 67\} + \alpha_{68}p_{ij}\mathbf{1}\{Age = 68\} + \alpha_{69}p_{ij}\mathbf{1}\{Age = 69\} + \alpha_{70}p_{ij}\mathbf{1}\{Age = 70\} + brand_j + \epsilon_{ij}$, $\epsilon_{ij} \sim iid\, Type\, 1$. The specification includes fixed effects for eight largest insurers. The estimates use separate cross-sectional parts of data sample that is used later to estimate the full choice model. The sample is further restricted to only include individuals that are 65-70 years old. The estimates show that in the later years of the program the price sensitivity of new and existing enrollees diverges in the direction that is consistent with the hypothesis of substantial switching costs - enrollees with incumbent plans appear significantly less price sensitive (and similarly so across different 66+ ages) than newly entering beneficiaries. Standard errors (not reported) are clustered at the regional level; p-values in square brackets - differences from the baseline of 65 year olds significant at <5% are marked in bold font. Reported are coefficients on premiums in the utility function and not marginal effects.

Table 1.6: Contract choice model specifications

	Non-IV (1)	IV (2)	Non-IV (3)	IV (4)
Annual premium, $100	-0.3911	-0.4464	-0.4148	-0.4984
	(0.0074)	(0.0111)	(0.0078)	(0.0110)
Deductible, $100	-1.3633	-1.3745	-1.2329	-1.2377
	(0.1307)	(0.1322)	(0.1266)	(0.1266)
sigma	0.5787	0.5904	0.4802	0.4704
	(0.0250)	(0.0253)	(0.0232)	(0.0238)
x Risk	0.0190	0.0139	0.0393	0.0384
	(0.0373)	(0.0377)	(0.0367)	(0.0367)
Initial Coverage Limit, $100	-0.1584	-0.1589	-0.1650	-0.1679
	(0.0234)	(0.0239)	(0.0262)	(0.0268)
sigma	0.0679	0.0731	0.0815	0.0866
	(0.0045)	(0.0046)	(0.0053)	(0.0053)
x Risk	0.0934	0.0938	0.1052	0.1087
	(0.0059)	(0.0061)	(0.0066)	(0.0068)
Partial coverage gap, 1/0	-1.8659	-1.7814	-2.0850	-2.1528
	(0.2990)	(0.2948)	(0.3358)	(0.3336)
sigma	0.8055	0.6929	1.2640	1.2380
	(0.0611)	(0.0728)	(0.0522)	(0.0542)
x Risk	1.0315	1.0339	1.0954	1.1261
	(0.0814)	(0.0810)	(0.0897)	(0.0895)
Default plan, 1/0	5.4487	5.6093	5.0675	5.0914
	(0.2533)	(0.2596)	(0.2584)	(0.2618)
x Risk	0.2291	0.2227	0.3589	0.3687
	(0.0643)	(0.0658)	(0.0655)	(0.0664)
Observations	2,435,171	2,435,171	2,435,171	2,435,171
Likelihood at convergence	-62,470	-62,379	-59,291	-59,224
Number of insurer fixed effects	3	3	10	10
Switching cost 75y.o. female, av. risk	$1,506	$1,330	$1,392	$1,164

Table 1.7: Detailed parameter estimates of the preferred choice model specification (#4 in Table 1.6)

			Interaction with demographics									
Plan characteristic	Estimate	s.e.	Risk score	s.e.	White	s.e.	Female	s.e	Age	s.e	ESRD	s.e.
Annual premium, $100	**-.4984**	0.011	-	-	-	-	-	-	-	-	-	-
Deductible, $100 - μ	**-1.238**	0.127	0.038	0.037	-0.023	0.057	0.013	0.025	**0.008**	0.002	-0.039	0.258
Deductible, $100 - σ	**0.4704**	0.024	-	-	-	-	-	-	-	-	-	-
ICL, $100 - μ	**-0.168**	0.027	**0.109**	0.007	**0.021**	0.012	0.003	0.005	0.0003	0.0003	**0.099**	0.036
ICL, $100 - σ	**0.087**	0.005	-	-	-	-	-	-	-	-	-	-
Plan has LIS, 1/0	**-1.093**	0.241	**-0.125**	0.072	-0.089	0.106	0.017	0.047	**0.026**	0.003	-0.683	0.563
Fixed co-pays, 1/0	**-1.538**	0.394	**0.589**	0.119	**-0.343**	0.188	0.003	0.076	**0.0197**	0.005	**-1.521**	0.789
Partial gap, 1/0 - μ	**-2.153**	0.334	**1.126**	0.089	0.105	0.150	-0.059	0.065	**0.016**	0.004	0.055	0.540
Partial gap, 1/0 - σ	**1.238**	0.054	-	-	-	-	-	-	-	-	-	-
Default plan, 1/0	**5.091**	0.262	**0.369**	0.066	**-0.815**	0.125	-0.067	0.046	**0.015**	0.003	-0.394	0.381
Insurer fixed effects	Yes (10)											
Log-likelihood	-59,224											
N total observations	2,435,171											
N unique individuals	12,769											
N choice situations	47,770											
Max N choices	66											
N of regions	34											

The MSL estimates are based on the control function IV specification of the model as described in the main text. To reduce the computational burden, the model was estimated on a sub-sample of individuals. This allows me to include panel observations for these individuals for all fours years of the data as well as to directly specify each contract in the individuals' choice set in each year without any aggregation. To construct the estimation sample, I restricted the panel sub-sample to include individuals marked with CMS 5% sample flag (which amounts to taking a 25% draw of the panel sub-sample, since the original data represented a 20% population sample of Medicare). I then took a 5% random draw of that data in a way that preserves the original panel structure. The resulting dataset is not different from the original panel sub-sample in a statistically meaningful way. The table reports estimates of utility parameters and not marginal effects. Coefficients significant at 10% level are marked in bold font.

Table 1.8: Choice model fit: summary statistics by contract type and insurer for enrollment and risk distribution moments

	Enrollment		Average ex-post drug spending		Average risk score	
	Observed	Model	Observed	Model	Observed	Model
Contracts of type 1	21.71%	20.71%	$1,580	$1,614	0.85	0.84
Contracts of type 2	65.93%	69.78%	$1,743	$1,819	0.88	0.89
Contracts of type 3	10.78%	8.88%	$2,846	$2,716	1.01	1.03
Contracts of type 4	1.58%	0.63%	$3,971	$2,658	1.04	1.08
Insurer A	29.65%	30.77%	$1,915	$1,939	0.92	0.92
Insurer B	27.10%	25.91%	$1,621	$1,614	0.86	0.85

The table compares three within-sample predicted and observed moments in the data: 1) Enrollment shares in different types of plans and in different insurer brands; 2) Average drug spending in different types of plans and in different insurer brands and 3) Average risk scores in different types of plans and in different insurer brands. The data is pooled over time and regions. To simplify the contract space, the comparison is made at the 4-type plan aggregation and at brand-level aggregation for the top 2 insurers. A more disaggregated fit of the model is illustrated in Figure 1-7. For the risk scores and drug bills, "predicted" measures refer to the sorting of the observed risks and expenditures as suggested by the simulation of the choice model.

Table 1.9: Basic descriptive evidence generated in the model: share of enrollees choosing the "default" option

Includes all plans	2007	2008	2009
1. Share observed in the baseline sample			
Probability of choosing default plan for 66+ y.o. enrollees	89.9 %	88.7 %	89.1 %
N	1,089,978	1,162,545	1,194,036
2. Share observed in the estimation sample			
Probability of choosing default plan for 66+ y.o. enrollees	89.9%	89.5%	89.6%
N	11,170	11,640	12,197
3. Share predicted in the estimation sample			
Probability of choosing default plan for 66+ y.o. enrollees	86.3%	84.8%	86.3%
N	11,170	11,640	12,197

This tables reports the simulation of the baseline descriptive evidence on the switching rates as documented in Table 1.3 using the estimated contract choice model.

Table 1.10: Counterfactual simulations measuring the interaction between adverse selection and switching costs under the observed regulatory regime

Year 2009 outcomes	Type 1 contracts	Type 2 contracts	Type 3 contracts
Enrollment share			
A. Baseline prediction, observed prices	18%	72%	9%
B. No switching cost; observed prices	23%	70%	7%
C. No switching cost; re-pricing	26%	64%	9%
Average risk (risk scores in USD)			
A. Baseline prediction, observed prices	$1,842	$1,926	$2,368
B. No switching cost; observed prices	$1,892	$1,930	$2,355
C. No switching cost; re-pricing	$1,907	$1,917	$2,321
Average ex-post drug spending			
A. Baseline prediction, observed prices	$1,741	$1,881	$2,924
B. No switching cost; observed prices	$1,881	$1,939	$2,325
C. No switching cost; re-pricing	$1,915	$1,914	$2,330
Weighted average premium			
A. Baseline prediction, observed price	$407	$439	$842
B. No switching cost; observed prices	$350	$412	$767
C. No switching cost; re-pricing	$313	$456	$742

The reports results of a simulation that includes the model's prediction for the baseline with switching costs as well as two counterfactuals without switching cost for one year in the program - 2009. Counterfactual simulation of the baseline takes premiums and contract defaults in 2009 as given. Counterfactual simulation in B takes premiums in 2009 as they were observed on the market and shuts down the switching cost channel in the utility function. Counterfactual simulation marked with C allows premiums to adjust to the new sorting of individuals when switching costs are not present. These counterfactual premiums are calculated using the pricing model discussed in the text. Since the model for premiums assumes that insurers adjust prices in accordance with lagged expenditures in plans, switching costs only change pricing-relevant sorting in 2008. Price simulation thus takes the premiums in 2006 and 2007 as given and re-calculates premiums for 2008 and 2009. The choices without switching costs in 2008 are then simulated using the new prices, which in turn affect the simulation of prices and subsequent choices in 2009.

Table 1.11: Counterfactual minimum standard policies with costly & costless switching

Year 2009 outcomes	Type 1 contract	Type 2 contract	Type 3 contract
		Enrollment shares	
A. With switching cost, observed prices			
Baseline predicted 2009 enrollment	18%	72%	9%
Set minimum standard in 2009 at 2006 level	19%	71%	9%
High deductible (set SDB deductible to $1,000)	10%	80%	10%
Minimum standard without donut hole in 2009	19%	72%	9%
B. No switching cost, observed prices			
Baseline predicted 2009 enrollment	23%	70%	7%
Set minimum standard in 2009 at 2006 level	25%	67%	7%
High deductible (set SDB deductible to $1,000)	12%	80%	8%
Minimum standard without donut hole in 2009	24%	70%	5%
C. No switching cost, re-pricing			
Baseline predicted 2009 enrollment	26%	64%	9%
Set minimum standard in 2009 at 2006 level	27%	64%	9%
High deductible (set SDB deductible to $1,000)	20%	72%	9%
Minimum standard without donut hole in 2009	40%	47%	13%
		Average risk	
A. With switching cost, observed prices			
Baseline predicted 2009 sorting	$1,842	$1,926	$2,368
Set minimum standard in 2009 at 2006 level	$1,838	$1,929	$2,366
High deductible (set SDB deductible to $1,000)	$1,854	$1,914	$2,369
Minimum standard without donut hole in 2009	$1,870	$1,939	$2,246
B. No switching cost, observed prices			
Baseline predicted 2009 sorting	$1,892	$1,930	$2,355
Set minimum standard in 2009 at 2006 level	$1,896	$1,929	$2,366
High deductible (set SDB deductible to $1,000)	$1,901	$1,925	$2,326
Minimum standard without donut hole in 2009	$1,940	$1,962	$1,884
C. No switching cost, re-pricing			
Baseline predicted 2009 sorting	$1,907	$1,917	$2,321
Set minimum standard in 2009 at 2006 level	$1,901	$1,923	$2,311
High deductible (set SDB deductible to $1,000)	$1,948	$1,910	$2,320
Minimum standard without donut hole in 2009	$1,965	$1,947	$1,935

1.7 Appendix

1.7.1 Conceptual framework

A stylized model of insurance contract choice below highlights the key economic channels connecting the changes in contract features, e.g. stemming from minimum standard regulation, with adverse selection and switching costs. Consider a mass of beneficiaries, each described by a pair of characteristics - the individual's risk type r, as well as risk preferences and other demographic or idiosyncratic factors that may affect the individual's preference for insurance together denoted with ϕ. For simplicity, assume that the individual faces a choice between two insurance contracts that differ only in their deductible. The more generous contract H has a zero deductible and a premium p_H, while the less generous contract L has a deductible $d > 0$ and a premium $p_L < p_H$. Assuming the separability of prices in the indirect utility function and letting $v(d, \phi, r)$ denote the valuation of a contract with deductible d by individual (ϕ, r), we arrive at a standard choice problem in a vertically differentiated goods environment. Individual (ϕ, r) chooses contract L if:

$$v(0, \phi, r) - v(d, \phi, r) < p_H - p_L$$

$$\triangle v(d, \phi, r) < p$$

where p denotes the relative price. Suppose that for any given level of the deductible, the valuation of an insurance contract is increasing in risk r, i.e. $\frac{\partial v(d,\phi,r)}{\partial r} > 0$ and preferences such as risk aversion, i.e. $\frac{\partial v(d,\phi,r)}{\partial \phi} > 0$, while the valuation is decreasing in the deductible for a given (ϕ, r), i.e. $\frac{\partial v(d,\phi,r)}{\partial d} < 0$. Suppose further that the valuation and prices are such that the "market is covered" in the sense that all individuals find it optimal to buy one of the insurance contracts rather than to remain uninsured.[35] Then, there exists an individual of type $(\hat{\phi},\hat{r})$ who is indifferent between the two contracts, i.e. $\triangle v(d, \hat{\phi},\hat{r}) = p$. The average risk that contract L expects to get after individuals choose between the two contracts is $E[r|\triangle v(d, \phi, r) < \triangle v(d, \hat{\phi},\hat{r})]$.

To introduce the policy of a minimum standard regulation in this model, let this policy be one-dimensional and only set the maximum allowed deductible $\bar{d}$. Assume further that the less generous contract sets its deductible d to always equal the maximum deductible set by the government: $d = \bar{d}$. The more generous contract, at the same time, always keeps zero deductible. This simplification implies that I am not modeling how insurers originally decide whether to offer the minimum standard or zero deductible, taking these decisions as given and stable from the policy perspective.

Now suppose the government changes its policy and increases the maximum allowed deductible from d to $d' > d > 0$ and nothing else changes. In particular, suppose for a moment that relative prices remain the same p. Individuals that were choosing contract L before, will switch to contract H under the new policy if now:

$$\triangle v(d', \phi, r) > p$$

The risk pool of switchers from the less to the more generous contract under the new policy but without price adjustment is: $E[r|\triangle v(d, \phi, r) < p$ and $\triangle v(d', \phi, r) > p]$. Whether this re-sorting results in higher or lower risk in contract L depends on whether the effect of risk on valuation grows faster at a higher deductible than the effect of non-risk preferences on valuation under a higher deductible. In other words, it depends on the relationship between $\frac{\partial^2 v(.)}{\partial r \partial d}$ and $\frac{\partial^2 v(.)}{\partial \phi \partial d}$.

[35] While this assumption is certainly restrictive and eliminates an important extensive margin on which the minimum standard may affect the market (Finkelstein, 2004), the empirical model in this paper focuses on the effects of the minimum standard on the intensive margin, across different levels of contract generosity, and thus I focus on this aspect of the question in this stylized model as well.

Now suppose that individuals face a switching cost γ. This cost may be heterogeneous across individuals and correlate both with individual preferences ϕ and risk type r. Let γ be a function of individual characteristics $\gamma(\phi, r)$. The switching cost could be both increasing or decreasing in risk and this relationship will influence whether switching costs increase or decrease the level of adverse selection conditional on the change in deductible. With the switching friction individuals that were choosing contract L before the policy change, will switch to contract H under the new policy if:

$$\triangle v(d', \phi, r) > p + \gamma(\phi, r)$$

The switching cost has the effect of reducing and tilting the set of beneficiaries that are indifferent between switching to H and staying in L. Whether relatively higher or lower risks tend to stay in contract L rather than change to H in the presence of switching cost will depend on the partial and cross-partial derivatives of the switching cost with respect to risk r and preferences ϕ. Allowing insurers to adjust prices to the new regulation and sorting patterns that are distorted by the switching costs produces theoretically ambiguous results that depend on the relationship between contract valuation and risk. For example, with a higher regulated deductible, the relative price will increase because a higher deductible mechanically reduces the liability of contract L. This, in turn tightens the switching constraint $\triangle v(d', \phi, r) > p' + \gamma(\phi, r) > p + \gamma(\phi, r)$, which can further decrease or increase the risk depending on the individual value function. Overall, the direction of change in sorting patterns induced by the change in the contract space are ambiguous if we allow for switching costs and allow insurers to adjust prices in response to changes in selection patterns. The effect that the regulation has on the allocation of risks across contracts will depend on the partial and cross-partial derivatives of the valuation and switching costs with respect to risk and preferences. The choice model in Section 1.4 estimates these inter-dependencies in Medicare Part D empirically and uses the estimates to simulate the role of switching costs in shaping the risk-sorting across contracts in response to market-driven and regulatory changes in contracts.

Table 1.12: Construction of the baseline sample

	2006		2007		2008		2009	
Full sample, N	9,999,538	*100%*	10,176,611	*100%*	10,369,814	*100%*	9,781,213	*100%*
Keep age 65 + and residence within 50 states	8,385,276	*84%*	8,511,573	*84%*	8,658,693	*83%*	8,066,696	*82%*
Drop if individuals died in the reference year	7,982,664	*80%*	8,111,023	*80%*	8,249,112	*80%*	7,714,002	*79%*
Drop if dual eligible any month	6,839,959	*68%*	6,952,339	*68%*	7,087,638	*68%*	6,637,418	*68%*
Medicare entitlement b/c of old age	6,412,259	*64%*	6,505,996	*64%*	6,619,029	*64%*	6,178,410	*63%*
Keep PDP enrollees [a]	1,797,409	*18%*	1,739,617	*17%*	1,800,364	*17%*	1,611,820	*16%*
Drop recipients of premium subsidies	1,551,253	*16%*	1,597,567	*16%*	1,668,923	*16%*	1,505,854	*15%*
Drop RDS and missing risk scores	1,221,252	*12%*	1,307,966	*13%*	1,356,861	*13%*	1,365,239	*14%*
Baseline sample	**1,221,252**	*12%*	**1,307,966**	*13%*	**1,356,861**	*13%*	**1,365,239**	*14%*
Panel sub-sample	871,818	*9%*	911,403	*9%*	954,494	*9%*	998,014	*10%*

The table shows the restrictions to the original sample of 20 % Medicare beneficiaries that were imposed to get to the baseline sample. The key restriction was to drop observations on individuals who didn't enroll in any Part D plan or enrolled in Part D through their managed care plan rather than through a stand-alone prescription drug plan (PDP). For years 2007-2009, I kept only individuals who were enrolled in a PDP for the whole year with the exception of the 65 year olds - this excludes those individuals who were allowed to join the plan outside of the open enrollment period because they e.g. changed their state of residence. In 2006, given the different special open enrollment period, many individuals were not enrolled for all 12 months and so I keep all individuals who initiated enrollment at some point during 2006 and didn't leave in subsequent months of 2006.

[a] Mainly drops those who did not enroll in Part D at all and those who enrolled in Medicare Advantage or other Part D coverage options.

Table 1.13: Pricing model used for the simulation of premiums in the counterfactual scenarios

$$E[Y_{jbt}|.] = \alpha_b + \delta_r + M'_{jbt-1}\beta + \gamma_1 Ded_{jbt} + \gamma_2 ICL_{jbt} + \gamma_3 1\{PartialGap\}_{jbt}$$

where j indexes region-specific plans, b indexes insurers (brands), r indexes 34 Part D regions, t indexes years

	(1)
	Annual premium, USD
Lagged mean spending	0.132***
	(0.00992)
Deductible amount, USD	-0.489***
	(0.0262)
ICL amount, USD	0.312***
	(0.0198)
Partial coverage in the gap, 1/0	293.9***
	(11.89)
Insurer FE	Yes
Region FE	Yes
N	2566
Mean Y	540.2
St. dev. Y	253.3
R-squared	0.802

Clustered standard errors in parentheses

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

The pricing regression is estimated on a dataset that records, for all prescription drug plans, the annual premium, the mean, the standard deviation and other moments of the lagged drug spending distribution in the plan (by plan enrollees in the baseline sample). The data also records the key financial characteristics of the plans - the deductible, the ICL and the gap coverage indicator of each plan in the program for years 2007-2009. For the cases where plans changed their ID over time due to mergers, I use Medicare plan cross-walk to match plans. The regression output doesn't report the coefficients on the set of fixed effects, as well as on the standard deviation, the kurtosis, the inter-quartile range, the 95th and 5th percentiles of the lagged distribution of realized expenditures.

Chapter 2

The Welfare Effects of Supply Side Regulations in Medicare Part D [1]

2.1 Introduction

The past decade has seen an unparalleled increase in the scope and magnitude of the private provision of government-subsidized social programs. The motivation for private provision is that the combination of supply-side competition and consumer choice would maximize consumer utility and keep costs low. One such program, the expansion of Medicare into prescription drug coverage through Part D in 2006, has generated a tremendous amount of policy interest and academic research. So far, the vast majority of the literature has been focused on demand-side questions.[2] At the same time, the supply side of Medicare Part D is vital to the operation of the market, as the program is centered around publicly-subsidized private delivery of prescription drug plans. In this paper we start closing this gap by examining the role of Medicare Part D's pricing regulations on supplier behavior.[3] We show that the mechanisms through which Medicare currently administers its subsidies on this market creates incentives for the suppliers to strategically manipulate the market in ways that reduce the overall welfare of the program. We investigate several alternative mechanisms for determining the premiums and subsidies on this market and analyze their implications for welfare in counterfactual simulations.

Medicare Part D is an elective program providing prescription drug insurance coverage to eligible

[1]**This Chapter is co-authored with Francesco Decarolis and Stephen Ryan.** Decarolis is grateful to the Sloan Foundation (grant 2011-5-23 ECON). We also gratefully acknowledge support from the NSF (SES-1357705)

[2]See Heiss, Mcfadden, and Winter (2010); Heiss, Leive, McFadden, and Winter (2013); Abaluck and Gruber (2011, 2013); Kesternich, Heiss, McFadden, and Winter (2013); Ketcham, Lucarelli, Miravete, and Roebuck (2012); Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2012); Einav, Finkelstein, and Schrimpf (2013)

[3]The literature on the supply side of Part D is still rather small. The studies that are closer to our project are Ericson (2013), that looks at how insurers increase premiums through time to exploit consumers inertia in plan choices, Duggan, Healy, and Scott Morton (2008); Duggan and Scott Morton (2010), that estimate the effect of Part D on drug prices, and Yin and Lakdawalla (2010), that analyzes how Part D enrollment affects private insurance markets.

populations. Consumers face subsidized insurance premiums and the subsidies vary across two groups: so-called regular enrollees and low-income enrollees. Regular enrollees make unrestricted choices from all plans offered in their region and pay a (partially) subsidized premium. In contrast, many low-income enrollees, who constitute 35 percent of all enrollees, are randomly assigned to eligible plans by CMS and pay nothing.[4] The level of subsidies for the regular enrollees are set as a non-linear function of bids that all insurers submit to Medicare annually. Critically, the same bids determine both the subsidy levels for the low-income enrollees and the plans' eligibility to enroll the randomly-assigned LIS enrollees. Consequently, there is key linkage between the two groups: the bidding process by which plans qualify to be eligible for low-income assignments also influences premiums for regular enrollees. Thus, these incentives distort both the public payments for low-income enrollees and the prices and choices of regular enrollees.

Our goal in this paper is to understand and disentangle market distortions from the process by which subsidies for the regular enrollees are set and the distortions imposed by the LIS-related random assignment. On the latter, Decarolis (2013) documents that there is considerable empirical evidence consistent with strategic behavior on the part of firms in manipulating the low-income subsidy. While this reduced-form evidence shows the scope for strategic manipulation of the low income subsidy, policymakers would ideally like to have an informed view of the costs and benefits of the current mechanism compared against other policies. Moreover, it is important to take into account the relationship between the LIS random assignment and the subsidy levels of regular enrollees. Therefore, in this paper we pose and estimate a structural model of supply and demand in the Medicare Part D market. Having a structural model of the market allows us to make predictions about outcomes in regulatory environments not observed in the data.

We are first interested in understanding the welfare consequences of the mechanism through which the subsidies to regular enrollees are determined. We analyze a set of partial and full equilibrium counterfactual scenarios that directly address this question. Currently, the subsidies for regular enrollees are determined as follows: first, each year insurers submit bids to CMS for each plan they offer across all of the Medicare regional markets they participate in. These bids, after being adjusted for risk, determine the payment firms earn when a regular enrollee chooses their plan. Consumers, however, do not see these bids directly; instead, they receive a (partial) subsidy determined through a somewhat convoluted mechanism. CMS takes the bids for all plans in the entire United States and computes a national average bid. Approximately one quarter of this national average is then called the base premium. The base premium determines the split of costs between the government and consumers: for all plans priced at or below the base premium, the consumer-facing premium is set to zero. For all other plans, the premium is set to be the difference between the firm's bid and the base premium. This process distorts both the absolute and relative prices that consumers face when choosing among Part D plans.

[4]Low income enrollees can opt out of the random assignment process and freely choose their plan, but this entails some costs. These enrollees can also re-enter the random assignment system under certain conditions. As of 2011, about two thirds of these enrollees had not opted out of the random assignment system.

Using our structural model of supply and demand, we characterize the welfare effects of this pricing mechanism. We start with three benchmarks. We first calculate welfare estimates from the observed prices and allocations. Our findings suggest that relative to the existing outside option, the current levels of subsidies in the stand-alone Prescription Drug Plans are generating a small return on government spending and negative nominal welfare. In the second benchmark, we re-simulate the current bid-averaging mechanism in our model. This exercise implicitly removes the distortions from the random LIS assignment. We find that consumer surplus almost doubles in this case relative to the observed allocation and prices. The final benchmark that we calculate is marginal cost pricing, where we completely remove the subsidy to regular enrollees and force the insurers to price at the estimated marginal costs. This counterfactual nominally increases welfare by 10 billlion USD; however, the model predicts almost zero enrollment in PDP plans - this result is not surprising in light of our earlier finding that the observed prices and allocations generate relatively low surplus for dollar. To interpret these results, it is important to keep in mind that the outside option in our model includes the financially attractive prescription drug coverage in Medicare Advantage plans rather than a "no insurance" outside option.

With these benchmarks in mind, we then proceed to investigate a menu of counterfactual subsidy-setting policies that CMS could implement in lieu of the current bid averaging process. The simplest scenario would be to provide fixed vouchers that could be used to buy a plan in the Part D market. A second option would be to use a uniform proportional discount on all plans' bids. A third policy would be to move closer to the idea of auction mechanisms and set subsidies equal to the lowest bid. Finally, we also consider all of these outcomes against the case where the direct subsidy was eliminated and the premium reflected the bid directly.

We simulate premiums and allocations across a menu of counterfactual equilibria. From these simulations, we conclude that proportional subsidies that are not tied to the bids may generate higher consumer surplus and higher payoffs than average-tied subsidies that alter relative prices. At the same time, similar payoff on public funds and higher levels of welfare, can be achieved with optimally chosen flat subsidies that are specified ex ante to the annual market. The more extreme simulations of 0% to 100% subsidies (or dollar equivalents of those) suggest that higher subsidies in this market have a perverse effect of increasing the incentives for the insurers to increase their bids in the program. While the goal of subsidy programs is to increase the affordability of insurance, in practice we find evidence consistent with the idea that in oligopolistic markets, the insurers with market power will have the incentives to inflate prices and pass them through to the price-insensitive government.

The remainder of the paper is organized as follows: Section 2.2 describes the institutional details of the Medicare Part D market and our sources of data. Section 2.3 introduces the theoretical model underpinning our analysis, while Section 2.4 describes our empirical application of that model to the data and our results. Section 2.5 discusses our counterfactual pricing mechanisms and presents our results. Section 2.6 concludes.

2.2 Institutional Environment and Data

2.2.1 Institutional Setting

Medicare is a public health insurance program for the elderly and disabled in the United States that covers over 50 million beneficiaries. Signed into law in 1965, the program aims to provide health insurance for a population which is generally characterized by high health expenses and low economic resources, and which historically had trouble finding and affording private health insurance coverage. Medicare costs the government about $500 billion annually and constitutes a large (16 percent in 2012) and growing share of the federal budget.

The Medicare program is administered by the Centers for Medicare and Medicaid Services (CMS), and consists of several pieces. Parts A and B cover hospital and outpatient services, respectively, under a fee-for-service model. Part C or Medicare Advantage, introduced in 1997, allows consumers to switch from fee-for-service to government-subsidized managed care administered by private insurers. In 2006, Congress expanded Medicare program to include prescription drug coverage via Medicare Part D as part of the Medicare Modernization Act of 2003. In 2012, approximately 32 million individuals benefited from Medicare Part D program and the Congressional Budget Office estimates that the government currently spends over $60 billion on Part D annually. This new part of the program is the institutional setting of our study.

Medicare Part D coverage is voluntary and enrollment is not automatic for the so-called "regular" beneficiaries, who do not receive extra support from the government. Beneficiaries eligible for low-income subsidies are instead automatically assigned to plans by CMS. The "regular" beneficiaries, who choose their own plans, have to actively enroll in one of more than 30 stand-alone Rx contracts offered in their state of residence during annual open enrollment period or when they first become eligible, e.g. turn 65. Once enrolled, these beneficiaries pay premiums on the order of $400-$500 a year, and in return insurers pay for prescription drug purchases subject to a deductible, co-payments or co-insurance, and coverage limits.

The exact structure of the Part D benefit in Stand-Alone Prescirption Drug Plans (henceforth PDPs) is quite unusual. Medicare Part D's contract space is driven by a minimum coverage requirement from CMS, known as Standard Defined Benefit (SDB). The insurers are required to provide coverage that gives at least the same actuarial value as the SDB. The SDB has an unusual design that features a relatively low deductible, flat co-insurance rate of 25% up to the initial coverage limit (ICL) and subsequent "donut hole", or coverage gap, that has a 100% co-insurance until the individual reaches the catastrophic coverage arm of the contract. Figure 2-1 illustrates what these features imply for a beneficiary. Consider an individual who in 2006 was in an SDB contract and purchased prescription drugs for a total of $3,000. Out of this amount, the individual would pay the deductible of $250, then 25% of the next $2,000 up to the initial coverage limit of $2,500, and then 100% of the next $750 in the gap, for a total out of pocket spending of $1,500. The remaining $1,500 would be paid by the plan.

A crucial feature of the institutional setting, which generates variation in contract characteris-

tics, is that insurers are allowed to adjust and/or top up the SDB contract design as long as their contracts cover at least the same share of average spending as the SDB. As a result, contracts offered by PDP insurers are highly multidimensional and vary on a variety of characteristics that differentiate them from the SDB minimum. Some of this differentiation is purely financial - contracts can change cost-sharing thresholds, co-pay and co-insurance levels, and may offer coverage in the "donut" hole. Other differentiating features are more related to the quality of the insurance provider. In our empirical specification of demand, we will be projecting all contracts available in the market into the same set of these pecuniary and non-pecuniary features.

The supply-side of the Part D program also has a unique, and controversial, design. Unlike the rest of Medicare, the drug insurance benefit is administered exclusively by private insurance companies. At the same time, the setting differs from more conventional private insurance markets in two key ways. First, the participating insurance companies are highly regulated, and continuously audited by Medicare. Second, consumers bear only a fraction of the cost in the program, as nearly 75 percent of insurer revenues from regular enrollees come from per capita Medicare subsidies. For individuals with low incomes, these subsidies go up to 100 percent. In this paper we focus on the distortions in pricing incentives generated by the complex regulatory structure of the subsidy mechanisms.

We focus on two features of the subsidy policies. First, to decide upon the division of the total per enrollee revenue between the consumer premium and the subsidy component, the government administers an annual "auction" mechanism. According to this mechanism, all insurers wanting to participate in the program in a given year submit bids for each plan they will be offering. Part D program is divided into 34 geographic markets and insurers are allowed to submit separate bids for the same plan in different regions. By statue, the bids are supposed to reflect how much revenue the insurer "needs" (including a profit margin and fixed cost allowances) to be able to offer the plan to an average risk beneficiary.[5]

Medicare takes the bids submitted by the insurers for each of their plans and channels them through a function that outputs which part of the bid is paid by consumers in premiums and which part is paid by Medicare as a subsidy. This function takes the bids of all plans nationwide, weights them by enrollment shares of the plans and takes the average. Roughly 75 percent of this average is the Medicare's subsidy portion. The remaining 25 percent of the national bid average, as well as the difference between the plan's bid and the national average is set as the consumer's premium. The per capita subsidy payment from Medicare is further adjusted by the risk score of each enrollee,

[5]There are several nuances buried in the set-up of the bidding procedure that are important for insurers' incentives and will enter the insurers' profit function in our empirical model. First, Medicare sets a minimum required actuarial benefit level that plans have to offer. Plans are allowed to offer more coverage ("enhance" the coverage), but that enhanced portion is not subsidized. Thus, when submitting their bids plans are supposed to only include the costs they expect to incur for the baseline actuarial portion of their benefit. The incremental premium for the enhanced coverage in the plans has to be directly passed on to the consumers. Second, while insurers submit one bid reflecting the costs required to cover an average risk in the program, the subsidy that the insurers receive for each individual in practice will be adjusted by the individual's expected risk. This adjustment is based on a risk-score index that reflects the differences in expected drug spending based on the individual's pre-existing medical conditions.

while the consumer premium may also include the additional payment for enhanced benefits if the plan offers them. This averaging procedure that we argue distorts relative prices, is the first piece of subsidy regulations that we focus on.

The second feature of the subsidy policies that we consider, concerns the role of low income beneficiaries (LIS) in the Part D program. Medicare utilizes the bids of the mechanism outlined above, to also determine the level of subsidies provided to the low income (LIS) population. For each geographic market, Medicare calculates the average consumer premium (without the enhanced coverage add-ons); the average is weighted by the lagged LIS enrollment in the plans. This average constitutes the subsidy amount that low-income beneficiaries receive, known as LIS benchmark or LIPSA. Most LIS beneficiaries do not in fact choose plans, but rather are randomly assigned by Medicare to qualifying plans in their regions. Qualifying plans are those that have premiums below the LIS benchmark and thus by definition zero premium for the LIS enrollees. Decarolis (2013) demonstrates that the way the LIS subsidy and enrollment are designed significantly distorts insurers' incentives and encourages gaming. In this paper we incorporate these incentives into the supply side model and consider the welfare implications of the interconnections between the LIS policy distortions and the bidding mechanism at large.

2.2.2 Data and Descriptive Facts

Our primary dataset combines a variety of aggregate plan-level statistics released annually by the centers of Medicare and Medicaid (CMS). Table 2.1 provides the basic summary statistics for our sample of Part D plans. We consider only stand-alone-prescription drug plans, which excludes prescription coverage that is bundled with Medicare Advantage health policies. Figure 2-2 documents the distribution of the outside option's share across markets in 2010. Since we are including both MA-PDs as well as other types of prescription drug coverage, together with the choice of no coverage, into the outside option, its share is quite high. The outside option share ranges from 50% to about 90% across markets.

In years 2007-2010, which are the focus of our empirical analyses, we observe a steady number of PDP plans totaling about 1,500 in each year.[6] This corresponds to 43–47 plans on average per market that individuals are choosing from. As Table 2.1 demonstrates, the supply-side is more concentrated than the raw counts of plans may suggest. We observe a total of around 50 insurer parent organizations operating in Part D in these years, with on average 19 separate organizations competing in each market. Figure 2-2 further shows that there is large heterogeneity in market shares attained by single plans both within and across markets. While many plans have market share close to zero, some plans cover as many as 20% of PDP enrollees within a market.

Table 2.1 shows that the average plan premiums for regular enrollees increased quite substantially in the time frame we are considering. The unweighted average premium went up from $423 per year in 2007 to $562 in 2010. As the box-plot diagram in Figure 2-3 illustrates this growth in

[6]This number slightly understates the total number of PDP available, as we had to drop several plans due to missing observations on some characteristics

premiums was accompanied with increased dispersion in plan premiums and in particular a higher number of very expensive plans. While Figure 2-3 suggests that the dispersion in premiums was relatively similar across different markets in the 2007 cross-section, the growth in premiums between 2007 and 2010 differed dramatically across regions.

The increase in the non-weighted average premium over the four years ranges from 13 percent in New Mexico to 61 percent in the California market. Part of the explanation for the different development of premiums could be the differences in market power exercised by insurers in different markets. Figure 2-4 documents a stark downward slope between the level of premiums and the number of competing parent organizations in a market. A different part of the explanation could be the differential effect of policy-design distortions across markets. Indeed, in our structural model of the supply side we will be able to disentangle some of these mechanisms and their heterogeneous effects across geographic regions.

In addition to the aggregate plan-level statistics, we use two complementary data sources. First, for the estimation of one of our demand model specifications, we also incorporate several moments constructed from the Medicare Part D's micro-level administrative data. These data include detailed claim-level information for a 5% random sample of all Medicare beneficiaries. From these data, we construct average by-market enrollment into plans with coverage in the gap separately for individuals older and younger than 75.

Finally, we amend the aggregate data with individual-level data from the Integrated Public Use Microdata Series (IPUMS) for years 2006-2010. These data allow us to construct the total potential enrollment in Medicare Part D, split by whether individuals are eligible for low-income subsidies. Further, we also sample age and income statistics from the survey that we utilize in the estimation of individual heterogeneity parameters in the BLP demand specifications.

2.3 Model

To evaluate the welfare implications of the highlighted regulatory mechanisms, we propose an empirical model of demand and supply of insurance contracts in Part D. The model takes into account the key policies governing the Part D subsidy system. We start with a model of demand for insurance contracts that follows the random coefficient logit approach (BLP) with additional micro-moments constructed from the administrative data (Petrin 2002). We then move to a supply-side model that allows us to estimate the marginal costs of the insurers. As we discuss below in more detail, we adjust the standard supply-side approach to take into account the distortions generated by the subsidies.

2.3.1 Demand

We utilize the standard approach of the discrete choice literature to estimate demand. Recalling that individuals eligible for the low-income subsidies do not choose their plans, but are instead assigned to plans by CMS, we restrict our attention to the regular enrollees. Further, we focus on

how regular enrollees choose among stand-alone Prescription Drug Plans (PDPs), and let the choice of not enrolling into the Part D program or enrolling through a Medicare Advantage plan bundled with health coverage be the outside option. We posit that individuals select insurance contracts by choosing a combination of pecuniary and non-pecuniary plan characteristics that maximizes their utility. We take the characteristics-space approach and project all plans into the same set of characteristics. This approach allows us to make fewer assumptions about how individuals perceive the financial characteristics of plans, but also implies that we remain agnostic about the objective "correctness" of choices. The characteristics-based approach suits our goals, as we are interested in capturing the demand response to changes in prices induced by counterfactual subsidy policies.

Let the indirect utility function of a regular enrollee i who chooses plan j in market t be given by:

$$u_{ijt} = -\alpha_i(\max\{0, p_{jt} - F\}) + x_{jt}\beta_i + \xi_{jt} + \epsilon_{ijt}, \tag{2.1}$$

where p_{jt} is the plan premium, F is a monetary voucher from the government to be spent on a Part D plan (set at zero in the current regulatory regime), x_{jt} contains observable characteristics of plan j in market t, ξ_{jt} is the unobserved characteristic of the plan, and ϵ_{ijt} is a random utility shock, distributed as a Type I Extreme Value. In our estimation we define the market to be one of the 34 statutory Part D geographic regions in years 2007 to 2010, for a total of 136 well-defined markets. The observable characteristics of plans j in market t, x_{jt}, includes the annual deductible, a flag for whether the plan has coverage in the donut hole, whether the plan is enhanced, several generosity measures of drug formularies, and vintage of the plans that allows us to capture consumer inertia. We also include fixed effects for parent organizations that capture individuals' preferences for particular insurance companies and insurer-level quality characteristics of plans, such as pharmacy networks and quality of customer service.

Observed and unobserved consumer heterogeneity enters through random coefficients on price variable and the dummy for whether a plan offers any coverage in the donut hole. We allow the mean of the coefficients to vary with income and age that are sampled from the IPUMS survey data for the relevant population. The unobserved heterogeneity may include individuals' differences in risk and risk aversion, which theory suggests are important determinants of insurance demand. The coefficients are thus specified as:

$$\alpha_i = \alpha + \Pi_\alpha D_i + \sigma_\alpha \nu_i \tag{2.2}$$

$$\text{where } \nu \sim \mathcal{N}(0,1) \tag{2.3}$$

$$\beta_{igap} = \beta_{gap} + \Pi_{\beta_{gap}} D_i + \sigma_{\beta_{gap}} \nu_i \tag{2.4}$$

$$\text{where } \nu \sim \mathcal{N}(0,1) \tag{2.5}$$

where D_i is a 2×1 matrix of demographic variables - income and age - and Π_α and Π_β is each a matrix of 1×2 coefficients measuring how tastes change with demographics.

We complete the utility model by specifying the outside good of not choosing to enroll in a stand-alone Prescription Drug Plan for the eligible population. This utility is normalized to zero. As

described above, we define the market share of the outside option as the fraction of Part D-eligible consumers enrolled in MA-PD plans or not at all enrolled in Part D.[7]

To estimate this model we augment the Berry-Levinsohn-Pakes (1995) approach for aggregate data with micro-moments from the micro-level administrative data in the spirit of Petrin (2002).[8] The micro-moments help us explain the demand for coverage in the gap among Part D enrollees. In particular, we consider the fraction of individuals of age below or above 75 that buy some coverage in the gap. This generates two additional demand moments for each of 34 markets in every year. These additional moments help us account for the fact that older—and thus usually sicker and more expensive—enrollees are more likely to choose plans with gap coverage, all else equal.[9]

We utilize several standard instrumenting strategies to address price endogeneity. For our baseline specifications with random coefficients, we combine the BLP-style instruments and Hausman-style instruments. The former include the total number of plans in the market, enrollment-weighted average characteristics of other plans offered by the same firm, and the characteristics of plans offered by other firms. The Hausman-style instruments include the previous-year enrollment-weighted average of prices of plans offered in other regions in the same macro region and in the other macro-regions by the same company, where macro-regions are defined as three large geographic areas in the US. Hausman instruments are further interacted with the mean and standard deviation of the income and age distributions in each region-year combinations. We argue that Hausman-style instruments are especially appealing in this setting given the regulatory structure of the market. Insurers are allowed to offer the same contracts at different premiums in different regions. Instrumenting the price in one region with the prices of the same contract in other regions, allows us to isolate the variation in prices that is common across these contract due to, e.g., particular agreements of a given insurer with pharmaceutical producers, and is thus not correlated with market-specific unobserved quality due to, e.g., local marketing. Table 2.3 suggests a very strong

[7]The population not enrolled is calculated as the difference between the population over 65 and total enrollment in PDP and MA-PD plans. Although MA-PD markets could be considered disjoint from the markets for the stand-alone Prescription Drug Plans (PDP), we include these plans in the outside option, since the averaging bidding mechanism for pricing involves (usually substantially lower) bids by MA-PD plans. Including these plans into the outside option helps us account for the level effects that the presence of the MA-PD plans may have on all PDP bids. Although the outside good includes other private drug insurance as well as other more esoteric coverage arrangements, few elderly use these options.

[8]The model is not directly estimated in the administrative data due to commercial privacy restrictions that prohibit us from recovering the identify of insurers and their parent organizations in the micro-data. Since being able to know which contracts are offered by the same insurer is critical for our supply-side strategy, we primarily rely on the market-level data for our estimation

[9]Ideally, we would like to include more direct measures of risk besides age in this exercise. However, while such measures could be constructed in the micro-data, we cannot construct the same corresponding variable in the sampled demographics from the survey data. Thus, we chose age as a variable that is both consistent and easily measurable across the two data sources, and at the same time meaningfully correlates with expected spending and risk. We verified the appropriateness of relying on age, by checking the variation in the probability of choosing a plan with gap coverage by age above or below the 75 year old group. Further, reduced form regressions with market level data confirm that at regional level moments of age distributions are correlated with the share of people that chose plans with gap coverage.

first stage.

While we instrument for the premiums of plans, we assume that other characteristics of the contracts are exogenous in the short run. We motivate this assumption by observing that insurers appear to offer a very stable portfolio of contract types over time. For example, if an insurer offers a contract with some coverage in the gap at the beginning of the Part D market, this insurers is likely to continue offering a contract with some coverage in the gap. The amount of coverage may change (and indeed, seems to have fallen over time), but the dummy-measure that we are using of whether there is any coverage in the gap does not appear to respond to short-term demand shocks or be related to anything else about the insurer and its plans. Similarly, for the deductible, the insurers tend to either follow the standard deductible set by Medicare every year, or reduce the deductible all the way to zero. We thus consider it appropriate to assume that short-run demand volatility and unobserved characteristics of the plans conditional on insurer fixed effects, such as advertising, primarily generate the endogeneity concerns for premiums, but not for the other features of the plans.

2.3.2 Supply

Modeling the supply side in Medicare Part D market presents a considerable challenge, as the decision-making of the insurers is affected by a complex set of regulatory provisions. We start with a description of the key regulatory distortions and set-up a general profit function that can incorporate these distortions. We then discuss a strategy we use to arrive at a simpler formulation of pricing incentives that allows us to map the supply side of Medicare Part D into the more familiar supply-side models of the differentiated products literature. We view the simplification strategy that we pursue as a contribution of the paper that allows us to get some traction on the supply side of this market relative to the existing supply-side literature in insurance.

We begin with a description of the revenue and cost channels for a single stand-alone prescription drug plan (PDP) in Medicare Part D. Let firm K offer plan j in a given region-year. Assume the firm does not offer any other plans in this market or other markets. We assume that all characteristics of plan j are pre-determined and the only decision variable for this firm is which bid to submit to Medicare for plan j. Suppose the firm does not qualify to enroll LIS enrollees that are randomly assigned by Medicare, but is eligible to enroll those LIS enrollees that opt out of the random assignment and choose their plans. For each individual that the plan enrolls, the firm collects a premium from the enrollees, p, as well as a subsidy from the government, σ. The premiums collected are the same for all enrollees—if the firm enrolls LIS individuals, that same premium is paid by the government. The premium is a function of the bid that this firm submits, b_j, as well as the average of all other bids in the whole country, $\bar{b}$. The subsidy payment from the government is different for each enrollee, as it is risk-adjusted. The level of subsidy depends again on the average bid, $\bar{b}$, as well as the risk measure of the enrollee, which can include LIS status. In other words, the subsidy is a function of average bid, health risk and LIS status: $\sigma_i(\bar{b}, r_i, LIS)$. On the cost side, the ex-post costs of the plan differ for each individual and depend on the individual drug expenditures. Some

of the costs are mitigated by the government through catastrophic coverage reinsurance provisions, according to which the government directly pays about 80 percent of individual's drug spending for particularly high expenditures. The costs of the plan also depend on the non-price characteristics of the plan, denoted by ϕ. We let individual-level ex-post costs be the function of health as measured by risk scores, r, and ϕ: $c_i(r_i, \phi)$.

The final piece of the firm's ex-post profit are risk corridor transfers among insurers and the federal government that happen at the end of the year at the parent organization level. These symmetric risk corridors restrict the amount of realized profits and losses that the insurers are allowed to collect in Medicare Part D. Thus, the transfers can be both positive or negative. We denote the function which adjusts a firm's ex-post profits by Γ.

The ex-post profit for plan j is then:

$$\pi_j(b) = \Gamma\left(\sum_i p_j(\bar{b}, b_j) + \sigma_i(\bar{b}, r_i, LIS) - c_i(r_i, \phi)\right). \tag{2.6}$$

For each individual, the subsidy and the cost can be expressed as an individual-specific deviation from the average subsidy and average cost: $\sigma_i = \sigma + \tilde{\sigma}_i$ and $c_i = c + \tilde{c}_i$. Denote the individual-specific difference in the subsidy and cost as $\eta_i = \tilde{c}_i - \tilde{\sigma}_i$. This function allows to capture post-risk-adjustment adverse or advantageous selection from the point of view of the insurance plan. Given the empirical evidence in Polyakova (2013) on the selection patterns in Medicare Part D, η_i mostly depends on whether or not the plan offers coverage in the gap rather than the premium per se. We thus let it be a function of plan characteristics: $\eta_i(\phi)$. Using this new notation, we can then re-write the profit function above as:

$$\pi_j(b) = \Gamma\left(N(p_j(\bar{b}, b_j) + \sigma(\bar{b}) - c(\bar{r}, \phi)) + \left(\sum_i \eta_i(\phi)\right)\right). \tag{2.7}$$

Denoting $\sum_i \eta_i(\phi)$ with $H_j(\phi)$, we obtain a profit function that does not have individual-specific terms and can be written using the market share notation that is useful for the empirical analysis. Note importantly that the premium combined with the average level of subsidies is equal to the bid of the firm by construction, i.e. $p_j(\bar{b}, b_j) + \sigma(\bar{b}) = b_j$. We can then re-write the pre-risk corridor profit of single-plan firm J for enrolling actively choosing enrollees as:

$$\pi_j(b) = (b_j - c)s_{jt}(b_j, b_{-j}, \bar{b})M^R - H_j(\phi). \tag{2.8}$$

We now expand this expression to multi-plan insurance organizations as well as allow the insurers to qualify for Low-Income-Subsidy enrollment. The LIS subsidy part of enrollment looks similar to the regular enrollees, except that the share function is different for these individuals. The function is non-linear, as the share is equal to zero if the bids are such that the plan doesn't qualify to enroll

the LIS. The profit function for firm J offering a portfolio of $j \in J_t$ plans across markets $t \in T$ is:

$$\pi_J(b) = \sum_{t \in T} \Gamma \left(\sum_{j \in J_t} M_t^R s_{jt}^R(b)(b_{jt} - c^R) - H_{jt}^R(\phi) + \sum_{j \in J_t} M_t^{LI} s_{jt}^{LI}(b)(b_{jt} - c^{LI}) - H_{jt}^{LI}(\phi) \right), \quad (2.9)$$

where (ignoring type superscripts) M_t is the population in the region, $s_{jt}(b_j, b_{-j}, \bar{b})$ is the share of plan j given the vector of all bids and the bid-averaging rule that translates bids into premiums, b_{jt} is the firm's bid for plan j in market t, and c is the marginal cost. Firms maximize profits by choosing b for each market.

While we able to incorporate a lot of institutional features of the regulatory environment, Equation 2.9 is still somewhat more complex than the typical differentiated-products profit function due to how the share equation, $s_{jt}(b)$, is constructed. For regular enrollees, the share depends on the premiums that are related to bids in a non-linear fashion according to the bid averaging rules:

$$p_{jt} = \max\left\{0, b_{jt} - \gamma \bar{b}\right\}, \quad (2.10)$$

where $\bar{b}$ is the average bid of all plans in the entire US and γ is the share of the average bid covered by the federal subsidy. The share equation for eligible plans in the low-income segment of the market is even more complex. It is composed of three parts: random assignment of new low-income enrollees by CMS, rollover of existing enrollees in the plan, and random reassignment of low-income enrollees within the parent organization from the plans that became ineligible in the current period. For ineligible plans, the share equation is equal to zero.

In the differentiated-product literature, marginal costs in Equation 2.9 are typically obtained by inverting its first order conditions. However, this inversion is complicated in our setting by the discontinuity in the share equation for low-income enrollees as well as the dependence of the government subsidy on the bids of all plans. In order to maintain the Nash-Bertrand assumption, we selectively study parts of the market where this assumption should be a reasonable approximation to the insurers' behavior.

First, we avoid imposing the Bertrand-Nash assumption on insurers actively competing to win low-income enrollees. Second, given the large number of plans in the country and the small influence of smaller plans on the bid averaging mechanism that is weighted by enrollment, we assume that the averaging mechanism does not distort pricing decisions for smaller plans. These two criteria implemented in different ways, as discussed in more detail in the estimation results, leave us with a considerable set of what we call "non-distorted" plans. Our strategy is then to invert the marginal cost using the standard differentiated-literature approach for these plans, and then project the estimated marginal costs from this part of the market onto those firms that we believe are distorted by the LIS or bid-averaging strategies.

2.3.3 Welfare function

The total welfare in Medicare Part D PDP market is comprised of several pieces. First, we have the consumers, or potential consumers on this market: seniors and the non-senior Medicare-eligible groups. With the introduction of Medicare Part D PDP plans, these individuals gain an opportunity to purchase stand-along prescription drug insurance coverage on exchange-style markets. We expect that this additional source of prescription drug insurance generates substantial consumer surplus; as Town and Liu (2003) conclude in their estimates of welfare effects from the introduction of Medicare Advantage program, prescription drug insurance is extremely valuable for the Medicare population. The second piece of the welfare calculation is the producer surplus that we measure using the pre-risk-corridor version of the profit in Equation 2.9.

The last important piece of net welfare calculations, are government transfers to the insurance firms and Part D PDP beneficiaries through federal subsidies. In our calculations, we consider both the direct subsidies that are used to reduce the individual premiums as well as the indirect subsidies in form of re-insurance payments used to ameliorate the ex-post costs of by the plans. In our welfare calculations, we weigh the government spending with the shadow cost of public funds set at 30 cents per dollar of the federal spending. As the Part D program currently accounts for a significant amount in budgetary outlays, an important question is how that compares to the consumer and producer surplus generated by the program. The second question is whether the net welfare calculation could be improved by the adjustment of the subsidy mechanisms in ways that would improve efficiency. This is one of the key metrics by which we assess our counterfactuals.

We calculate two kinds of welfare quantities using our model. First, we report welfare levels generated by PDP plans relative to the outside option for the observed set of plans and their prices. In essence, this number allows us to assess the value of the PDP portion of the Part D program. These calculations make an important *ceteris paribus* assumption about prices and characteristics other parts of Part D, such as MA-PD, as well as other potential sources of drug coverage that are all included into the outside option in our specification.

Second, we calculate changes in welfare that can be achieved through different subsidy mechanisms. To measure the levels and changes in consumer surplus in the observed and counterfactual allocations, we use the measure of welfare similar to that derived in Williams (1977) and Small and Rosen (1981) for the discrete-choice random utility models with extreme value type 1 errors. The measure is adjusted to account for the presence of random coefficients on price. This measure of consumer surplus assumes no income-effects in the utility function. While we do let the coefficient on price depend on individuals' income, in practice we don't estimate a statistically significant income effect, and thus view this measure as appropriate for our setting.

The change in consumer surplus thus takes the following form:

$$\triangle ConsumerSurplus = M \int \frac{ln[\sum_{j=0}^{J^0} exp(V_{ij}^0)] - ln[\sum_{j=0}^{J^1} exp(V_{ij}^1)]}{\alpha_i} \times dP(\sigma\nu)dP(D) \qquad (2.11)$$

In this expression, V_{ij}^0 is the representative utility for person i from contract j at the observed prices of the contracts. V_{ij}^1 is the representative utility for person i from contract j at the counterfactual prices of contracts. M is the total mass of individuals on the Medicare Part D market. The representative utilities are computed using the demand specification:

$$V_{ij} = -\alpha_i p_j + x_j \beta_i + \xi_j \tag{2.12}$$

To incorporate random coefficients, we simulate the integral in the expression above by making random draws from the estimated distribution of marginal utilities of income. The term $dP(D)$ accounts for the fact that we use market-level data in this estimation and thus do not observe the demographics of the individuals, which then have to drawn in the simulations. We use the same draws of the demographics as we used in the BLP estimation routine.

2.4 Empirical Results

The structural model above allows us to conduct predictions of market outcomes in economic environments not observed in the data. In Medicare Part D, several policy experiments naturally come to mind to help advance our understanding of the individual and collective effects on equilibrium outcomes that arise from regulatory choices on the supply side. Our counterfactual experiments focus on how the subsidy for regular enrollees is determined. Before turning to the analyses of counterfactual allocations and the associated welfare, we report the estimation results of the demand and supply models.

2.4.1 Demand

We present several specifications of demand estimates leading to our preferred baseline. We start with two specifications without random coefficients. The first one is the logit model estimated via OLS using the transformation as in Berry (1994). The linearity of this model allows us to easily test our instrumental variables in a 2SLS version of this specification before we go to the more involved estimation with random coefficients.

The results of the OLS regression of plans' market shares net of the outside good share on premiums and other characteristics of plans are reported in the first column of Table 2.2.[10] Despite

[10]For all demand estimates we use market-level data on all stand-alone prescription drug plans PDPs that were offered in years 2007-2010 and for which the data on pharmacy networks and drugs were available. The latter restriction excludes only a very small fraction of plans. We focus on the PDP plans only, as prescription drug plans that are part of the Medicare Advantage package are bundled with inpatient and outpatient insurance and thus it is hard to interpret demand for prescription drug insurance in that setting. We exclude the observations from year 2006, which was the first year of the program - in this year, the CMS did not have plan weights to calculate subsidies and all plans were weighted equally, thus changing the regulatory environment on which we focus. We stop out analysis at 2010 for similar reasons. First, there were major changes in the structure of plans in 2010, as the government started closing the gap in coverage. Second, CMS introduced a "meaningful difference" policy that required starker differentiation of plans and resulted in a significant number of plan exits or consolidations.

the general limitation of the logit model in producing reasonable substitution patterns, it gives us a simple way to establish that the variables included into the utility specification act as expected. In particular, we estimate a negative coefficient on the premiums and the deductible, suggesting that consumers dislike higher premiums and higher deductibles. We also estimate positive coefficients on the generosity features of the plans. As we would expect, beneficiaries like plans that offer coverage in the gap, cover more of common drugs and include more pharmacies as their in-network providers. We also note an economically and statistically significant positive coefficient on the vintage of plans, suggesting that plans that entered earlier in the program were able to capture a larger beneficiary pool.

In the second column of Table 2.2 we report the estimates of the 2SLS version of the Berry (1994) logit. Instrumenting has a large effect on the price coefficient, nearly doubling its magnitude. The estimates of the marginal utilities from other features also adjust slightly, retaining the intuitive signs. Table 2.3 reports the first stage and the reduced form for the IV exercise. We use a vector of instruments that combines the BLP and Hausman - style IV. The BLP instruments include several summary statistics based on the characteristics of plans that are competing with plan j. The Hausman instruments measure prices charged for the "same" plan in other geographic markets. The instrument picks up common cost-shocks for a set of plans offered by the same insurer, which may, for example, have especially favorable discount agreements with some pharmaceutical companies. The idea of the Hausman instrument is very appealing in the current setting due to the regulatory structure of the market. We assume that demand in different geographic markets, as defined by Medicare statue, is sufficiently disjoint to support the validity of the Hausman instrument together with the BLP-style IV.

We next proceed to the full BLP model that introduces "observed" and unobserved heterogeneity in the individual valuation of contract characteristics. The output of the standard BLP model is reported in the third column of Table 2.2. We find greater sensitivity to price in the specification that allows for unobserved heterogeneity and we also find some meaningful dispersion in the price coefficient. BLP estimation also produces some substantial adjustments in the magnitude of the coefficients for the other contract features. While all the signs are still as we would expect, we now find greater sensitivity to deductible levels, number of covered common drugs and in-network pharmacies. We do not find meaningful heterogeneity in preferences along the income dimension.

In our next specification, reported in column four of Table 2.2, we add a set of micro-moments generated from administrative micro-data on Medicare Part D to our BLP specification. Estimates with and without micro moments differ in an important and intuitive way. In our estimates without micro-moments, we systematically under-predict the share of individuals choosing the plans with gap coverage. Once we include micro moments, the estimated willingness to pay for having coverage in the gap increases substantially and we are able to closely match the share of people enrolling in plans with gap coverage. This result is of course not surprising, as we explicitly add the gap coverage dimension to the objective function through the micro-moments. The coefficient on price goes down by a bit in this specification, although it still remains almost three times higher than

the IV Logit. We take the IV Logit specification as our baseline estimates for the counterfactual analyses. Expanding the counterfactual analyses to the BLP model with micro-moments will be the subject of future work.

2.4.2 Supply

The key step in the supply-side estimation is the recovery of plan-level marginal costs that will enable us to simulate counterfactual prices under different regulatory scenarios of subsidy mechanisms. Unlike the standard differentiated product settings, our environment presents several challenges in way of profit function inversions. First, in general, insurance plans will not have constant marginal costs. Moreover, marginal costs will be a function of premiums and other characteristics of the plans, as these characteristics screen individuals of different expected risks. Second, the presence of subsidies that are determined through the mechanism that averages bids from all plans, potentially implies additional deviations from the standard Bertrand-Nash competition concepts. Finally, the presence of the low income subsidy market with its random assignment of individuals only to qualifying plans implies a discontinuity in the profit function.

Therefore, as discussed in the set-up of the supply-side, in order to proceed with the estimation of marginal costs, we make several important assumptions. First, we assume that the multitude of risk-adjustment and reinsurance mechanisms implemented in Medicare Part D imply that insurers de facto face constant expected marginal costs. Second, given the large number of plans in the country and the small influence of smaller plans on the bid averaging mechanism that is weighted by enrollment, we assume that the mechanism does not distort pricing decisions for smaller plans. Third, and similarly in spirit, we select a subset of plans that were plausibly not distorted by LIS gaming.

In essence, the idea is to select a group of plans for which we find the Bertrand-Nash assumption acceptable for describing the pricing behavior of the insurers. We construct two groups of such plans using different restrictions on the sample. In the first approach, we argue that an insurer that in a certain region never had any plan with a premium low enough to qualify for low-income enrollees is not pricing its plans to get the randomly assigned low-income enrollees. In the second approach, we select plans of insurers that within a given market (year-region) were not eligible to enroll randomly assigned LIS individuals into any of is plans. Although the assumption that these two groups of plans are "non-manipulating" appears reasonable, we may be worried that these plans are not comparable to plans qualifying for low-income enrollees. However, given the substantial variation in the low-income subsidy across regions, there are many insurers who never qualified for low-income enrollees in at least one region. This variation is mostly due to the different penetration of Medicare Advantage: where in 2006 enrollment in Medicare C was high MA-PD received a high weight in the calculation of the low-income subsidy and, since their premium is typically close to zero, they induced a small low-income subsidy.[11] Since the plans offered by the same insurer across

[11]The variation in the total weight assigned to MA-PD in 2006 is substantial ranging from almost 60 percent in Arizona and Nevada to less than 4 percent in Mississippi and Maine.

different regions are remarkably similar, the marginal cost estimates of the "non-distorted" plans through the inversion of the first-order condition, can be used to predict the cost of similar plans in other regions for which we could not directly apply the inversion approach.

Using the marginal cost estimates from the two groups of non-manipulating plans, we proceed to the next step of relating the estimated marginal costs to the observed characteristics of plans. In practice, we estimate the following hedonic-style linear regression:

$$mc_{jt} = X_{jt}\beta + \delta_t + \gamma_j + \epsilon_{jt}. \tag{2.13}$$

Table 2-5 reports two specifications of this regression using different sets of "non-distorted" plans. As expected, we note that the most important determinants of marginal costs for the insurance plans are the plans' key financial characteristics, such as deductibles and coverage limits. We use these estimates of how plan characteristics translate into marginal costs, to predict marginal costs for all plans that we did not include in the inversion procedure. This exercise hinges on the assumption that all plans have a similar "production function." In other words, we assume that the plans that manipulate the LIS threshold manipulate the premiums but do not have different marginal costs conditional on a set of non-price characteristics. This appears reasonable, as the main source of costs in the insurance market is determined by individual risk spending; therefore, it is conceivable to assume that plans with the same financial characteristics will have similar costs.

Figure 2-5 plots the distribution of predicted marginal costs and compares it to the estimated distribution via the inversion procedure. The figure displays the estimated annual marginal costs for the subset of plans that belong to insures that didn't qualify for random LIS assignment with any of its plans in a given market. We estimate substantial heterogeneity in the marginal costs. Observing heterogeneity in the estimates is important, as it indicates that our hedonic-style regression captures the key drivers for the differences in marginal costs—in other words, we don't get the predictions to be too centered around the mean, which would suggest that we were not capturing enough heterogeneity in marginal costs. The distributions appear broadly similar, which is not surprising, as the manipulating plans do not really differ from other plans in the key features of their contracts. We note that the distribution for the plans whose marginal cost we predict rather than estimate is visibly shifted to the left. This is an expected result, as the "distorted" plans that qualify to enroll LIS individuals should be cheaper and have lower costs.

2.4.3 Welfare

Using the demand and supply estimates, we compute consumer surplus, producer profits and government transfers for the observed allocation in the PDP part of the program. We restrict our calculations to regular enrollees that are not eligible for the low-income subsidy. Expanding the analysis to include LIS enrollees is an area for future work. The results of welfare calculations are reported in Table 2.4.

We estimate that at the observed prices, almost 9 million beneficiaries that are enrolled in PDP

plans attain an annual consumer surplus of 1.35 billion USD. This consumer surplus is estimated relative to the outside option, as the utility model is inherently ordinal. Producer profit estimated without the multitude of re-insurance transfers from the government, is estimated to be close to 0.9 billion USD.

On the cost side, we calculate that the government spends about 6 billion USD on premium subsidies for regular enrollees - equal to about 680 USD/year for each individual. Moreover, using CMS data on average non-premium level of subsidies, we calculate that the total amount of this subsidy is on the order of 3.5 billion USD. Multiplying the government subsidy by 1.3 to account for a 30 cents a dollar cost of public funds, we arrive at the lower bound of net welfare calculation for the regular enrollees to be about negative 10 billion USD. In other words, we estimate that a dollar of public funds generates at least 18 cents of welfare in the PDP program, relative to the very attractive outside options. These results present lower bounds, as we assume that the government would not spend anything on the individuals that exit the inside option as defined in the demand system. In practice, if individuals switch to the parts of the outside option that also carry some government spending, we may want to take into account that reduced opportunity cost of government funds. We will pursue this line of investigation in future research.

2.5 Counterfactuals

2.5.1 Partial equilibrium analysis

We first consider channels through which the subsidy mechanisms affect the allocation of individuals among plans conditional on the observed supply-side pricing decisions. In other words, we take the bids submitted by insurers for each plan in each region-year as given, and ask how beneficiary premiums and choices would change under different ways of subsidizing this market. In the partial equilibrium exercises we keep the LIS-related policies unaltered. We also keep the outside option that includes non-PDP parts of Part D market, such as MA-PD, unchanged.

The partial equilibrium counterfactuals rely only on the demand-side estimation and the changes across models do not depend on any of our assumptions about the insurer strategies. The observed subsidy mechanism used by CMS computes the subsidy by taking a percentage (about 68 percent) of the average bid of all Medicare Part D plan. CMS applies this subsidy to all plans in all regions of the country. In other words, the subsidy is not proportional to the generosity or price of the contract. We consider three alternatives to this mechanism. First, we ask how allocations and welfare would change if individuals were required to pay the full bid of the insurers. Second, we consider a mechanism closer to auctions that sets the subsidy to be equal to the lowest bid in a given market. Finally, we consider a proportional mechanism, where the subsidy is set as a fixed percentage of each plan's bid. For each of these mechanisms, we report changes in allocations, prices, producer profits, government budgetary expenditures, and consumer surplus.

Enrollees face the full bid without subsidy We start with a counterfactual that evaluates the consequences of removing all subsidies from consumer premiums among regular enrollees. In this counterfactual, we let regular enrollees directly face the observed bids of the insurers. In other words, we let:

$$p_{jt} = b_{jt}. \tag{2.14}$$

While it is clear that removing the subsidy will decrease individual welfare and enrollment, calculating the magnitude of this decrease allows us to compare it to the current government spending on the subsidy. There are several channels that determine the magnitude of the decrease in consumer surplus. First, all plans are now going to be substantially more expensive. Since we allow for an outside option and its implicit price doesn't change, the increase in the general level of prices in the inside option, is expected to induce some substitution on the extensive margin to the outside option. Moreover, due to the original subsidy mechanism, even in this partial equilibrium framework, removing the subsidy will change the relative prices of insurance plans and thus cause some substitution within the inside good.

We find that conditional on the *ceteris paribus* outside option that is quite attractive, substitution on the extensive margin is extremely important. Column two in Table 2.4 records that imposing the full bid on the consumers, is sufficient to induce almost all individuals to forego purchasing any stand-alone prescription drug insurance plan and substitute to the outside option. We estimate the the enrollment share in PDPs drops from the observed 28% to less than 0.5%. This may not seem surprising, as the average premim that beneficiaries face increases from on average 502 USD a year to 1,183 USD a year. It is again important to emphasize that the outside option here contains other options for prescription drug coverage, so individuals are not necessarily (and in fact rather unlikely) substituting to no prescription drug coverage at all.

Given the dramatic predicted drop in enrollment, it is not surprising that we estimate that consumer surplus and insurer profits drop to about 3 million USD combined. Interestingly, since now the government does not subsidize premiums (and we keep the non-premium subsidy), the value generated on the dollar of government funds almost triples from the baseline of the observed allocation. This total figure masks significant heterogeneity in the effect of the increase of premiums on consumer surplus across regions. Figure 2-7 records the distribution of average drop in per person consumer surplus across different markets in 2010.

Flat subsidy equal to the lowest bid In the next counterfactual we allow the government to subsidize consumers by providing them with a fixed subsidy for Part D plans. We let the subsidy amount equal the lowest observed bid in the data in each market. Picking the lowest bid within a market rather than over all plans in the country, allows accounting for permanent differences in the levels of prescription drug spending across different geographic areas. In other words, we let:

$$p_{jt} = b_{jt} - \min\{b_{jt}\}. \tag{2.15}$$

This counterfactual differs from the current policy on several dimensions. First, it changes the subsidy level for the consumers differentially across plans. Consequently, we expect substitution both within different PDP plans as well as on the extensive margin towards the outside option. Second, from the insurers' perspective, tying the subsidy to the lowest bid removes some of the potential gaming incentives. Moreover, it brings the system closer to a real auction setting, where the "winning" bid is rewarded by having the lowest (in this case zero) consumer premium. We come back to the effect of this mechanism on insurer incentives in the full equilibrium counterfactual, while here we evaluate whether setting the subsidy to the lowest bid would decrease the federal spending by more than consumer surplus drops.

We conduct the counterfactual on the last year of our data - 2010. We observe that recalculating the subsidy to be equal to the minimum bid, would decrease the premium levels faced by the enrollees by about a half. The general decrease in the premium levels accounts for a large enrollment effect - we estimate that enrollment increases for the observed 28% to 55%, or from about 9 million to about 18 million beneficiaries. This change in the subsidy also generates large changes in relative premiums. We estimate a total of a three-fold increase in consumer surplus (even though enrollment and less than doubles). Insurers profits from regular enrollees also go up to 1.6 billion USD.

At the same time, government expenditures increase quite dramatically. As column three in Table 2.4 records, the total PDP premium subsidies for regular enrollees increase from 6 billion USD to about 15 billion USD. Naturally, a large part of this effect is drive by the sheer enrollment effect - and it is important to note that we are not accounting that this is increase could be offset by a decrease in government spending on the outside option (i.e. on MA-PD subsidies). With this caveat in mind, we calculate that the total welfare relative to the outside option drop by 14 billion USD. At the same time, the surplus per dollar calculation, if anything, increases from 18 cents to 19 cents per dollar.

Proportional subsidy As the last counterfactual in the partial equilibrium analysis, we consider a subsidy that is proportional to the insurers' bids. We consider an array of different possible subsidy levels. We let the premiums range from 25 to 50 percent of the observed bids. A consumer premium equal to 25 percent of the bid is close to the goals of the current system, which claims to subsidize the individuals at about 75 percent. In practice, however, the observed subsidies are equal to about 65 percent of the average bid (in other words, the premium is equal to about 35%), so we test both levels in our counterfactual experiments.

We set the prices to equal x percent of the observed bid, with 25 percent being our baseline specification that is reported in the fourth column of the results Table 2.4.

$$p_{jt} = x * b_{jt}. \tag{2.16}$$

Even in the partial equilibrium framework, the welfare effects of changing to this subsidy mechanism are ambiguous, as going from a fixed amount to a proportional subsidy substantially changes relative prices. A proportional subsidy makes more generous plans relatively cheaper and thus

encourages individuals that value more generous coverage to substitute to these plans. At the same time, the overall level of prices may increase if 25% of the cheapest plan is going to be higher than the observed flat-rate subsidy levels. This would imply that the cheapest possible plan will be more expensive than under the observed regime and we will see substitution on the extensive margin away from Part D PDP coverage.

We indeed find evidence of important substitution patterns, with the allocation of individuals across plans changing dramatically relative to the observed allocation. With the premiums set at 25% of the observed bids, the average premium paid decreases from 502 USD a year to about 390 USD a year. At the same time consumer surplus increases almost five-fold to 6 billion USD. While the premium decrease is relatively modest, government spending goes up from 6 billion USD to 25 billion. With the enrollment up to 22 million beneficiaries, this still implies that the average government subsidy on premiums per beneficiary almost doubles. These observations all imply that within PDP plans, beneficiaries substitute to more expensive plans that are now relatively cheaper. Moreover, there is a lot of substitution to the inside option of PDP plans from individuals who were not in this part of the program before. While the total government spending goes up dramatically relatively to the counterfactual, where the subsidy was set equal to the minimum bid, the efficiency per dollar remains the same.

Figure 2-6 plots the levels of estimated consumer surplus at different levels of proportional subsidy. We find that the setting individual premiums to 47 percent of observed bids, or equivalently, setting the subsidy at 53 percent of the observed bid, results in the same consumer surplus as the observed allocation. This graph demonstrates the non-linear effect of the proportional subsidy on consumer surplus and suggests that we can in principle calculate the optimal proportion of the subsidy. Similar exercise can be conducted for fixed amount subsidies to determine the overall optimal. We discuss more of how such exercises may affect the supply-side pricing in the next section.

2.5.2 Full equilibrium analysis with regular enrollees

Full equilibrium counterfactuals allow the insurers to respond to the changes in the subsidy mechanisms together with the corresponding changes in individual demands. We conduct counterfactuals that try to understand the role of the bidding mechanism for pricing decisions in isolation from distortions induced by the random assignment of the low-income enrollees. In this set of counterfactuals, we essentially repeat the partial equilibrium exercises, but allow the insurers to change their bids, assuming that they only face regular enrollees. Table 2.5 breaks out the full equilibrium counterfactual changes in consumer surplus, producer profit, and government transfers between premium subsidies and other aspects of the PDP part of the Medicare Part D program.

Current subsidy mechanism without LIS distortion We start the full equilibrium analysis by calculating prices and allocations that our model would predict for the currently used subsidy mechanism. The key in this calculation is that we remove the distortion of the randomly assigned

LIS beneficiaries. Consequently, the results of this counterfactual allow us to assess the impact of the LIS distortion on prices. Moreover, this counterfactual generates a benchmark to which we can compare other subsidy mechanisms, keeping constant the separation of the LIS random assignment across the different subsidy mechanisms for the regular beneficiaries.

Specifically, in this counterfactual, we set the subsidy to be equal to 68 percent of the weighted average bids:

$$p_{jt} = b_{jt} - 0.68 * Average(b_{jt}). \tag{2.17}$$

Relative to the allocation observed in the data, we simulate lower consumer premiums (398 USD relative to 502 USD observed). Consequently, consumer surplus increases - doubling from 1.35 billion USD to 2.7 billion. Insurers' pre-reconcilliation profit slightly more than doubles. In response to lower premiums, beneficiary enrollment increases from about 9 million to 12 million individuals. While the nominal welfare calculation decreases by billion USD, the surplus generated per dollar of public funds increases from 18 to 23 cents a dollar.

We use the results of this counterfactual to isolate the effects of changing the subsidy mechanism from the current averaging method to other subsidy rules.

Enrollees face the full bid without subsidy, insurers are allowed to adjust bids We start with the simplest counterfactual, in which we remove any subsidies from the PDP market:

$$p_{jt} = b_{jt}. \tag{2.18}$$

Unlike in the partial equilibrium case, we now allow the insurers to re-price their plans. The thought experiment here is the following. Insurers now that with subsidies, the demand elasticity and enrollment depends on the plan premiums and not on bids, while the profits depend on bids. Since a fraction of the bid is paid by the government under the current mechanism, insurers may have an incentive to increase the bids, as only a part of this increase is passed-through on to consumers, and the rest is observed by the government that in our framework is price-inelastic. Thus, we may expect that bids would decrease if the subsidy is removed. At the same time, removing the subsidy changes the relative attractiveness of plans, which may have a countervailing effects on some bids. Lastly, in this counterfactual we are implicitly removing the distortion that came from the presence of randomly assigned LIS population. As now lower bids are not "rewarded" by these additional enrollees, it is possible that some insurers may increase their bids. In short, the effect of removing the subsidy on bids is ex ante ambiguous.

Column 4 in the lower panel of Table 2.5 reports the results of a counterfactual that is very close to the full removal of subsidies. This column reports the result of setting the premiums at 95% of the bids. In this case, we observe that bids do not change much from the observed levels. If compared to the partial equilibrium counterfactual, where individuals paid the full bid, the average premium is essentially the same, at about 1,200 USD a year. At the same time, the profitability of

the insurers decreases, from about 0.9 billion USD at the observed prices/subsidies to 0.02 billion, in line with the dramatic drop in enrollment to less than 0.5%.

Consequently, similarly to the results in the partial equilibrium calculations, we conclude that removing the subsidy nominally increases the efficiency of the program. Total welfare goes up by 10 billion USD, while the efficiency per federal dollar increases threefold to 56 cents a dollar. At the same time, this increase in welfare is driven by insurer profits and by the fact that beneficiaries almost completely substitute out of the market to the outside option of drug coverage.

Flat subsidy, insurers are allowed to adjust bids In this counterfactual we experiment with several different levels of flat vouchers that are applied to insurers' bids. We assume that insurers know the levels of the flat subsidy in advance and adjust their bids accordingly. In calculating the new equilibrium bids, we check corner solutions, where the insurers may decide to bid exactly at the subsidy level. The premiums are set equal to:

$$p_{jt} = b_{jt} - F. \tag{2.19}$$

We report results for two flat subsidy levels: of 676 USD a year, which is effectively the level of premium subsidies observed in the data. We also experiment with roughly doubling the subsidy to 1340 USD a year. In interpreting the first of these counterfactuals, it is important to emphasize that even though the subsidy is nominally the same, the mechanism is very different. While in the observed mechanism, the subsidy is determined after the bidding process as a fraction of the average bid, here we set the subsidy ex ante and it does not depend on the submitted bids.

In setting the flat subsidy equal to the level observed in the data, we find that the allocation and welfare change substantially relative to the observed allocation. Not surprisingly, we find government spending to be roughly similar to the observed levels, which is natural as we are setting individual premium subsidy to be the same as in the data. The level of prices increases, consequently both enrollment and consumer surplus decrease somewhat. Insurers' profits, on the other hand, increase by about 10%. Overall, welfare goes up by 3 billion USD and the efficiency generated by a dollar of spending increases to 25 cents.

A more extreme counterfactual is reported in column four of the lower panel in Table 2.5, where we increase the flat subsidy to 1,340 USD a year, which is double of the observed levels. This increase is sufficient to decrease increase the enrollment into the PDP part of the program to 88% or 27 million beneficiaries. In other words, with this level of subsidy, PDP plans are substantially more attractive than any outside option. This is not surprising, as the average premium with this level of subsidy drops to about 80 USD a year. Consumer surplus increases almost ten-fold to more than 11 billion USD. Producer profit increases substantially as well. At the same time, government spending increases much more dramatically, driving the welfare levels down by more than 35 billion USD. Despite this substantial drop in welfare levels, we still predict an increase in per dollar efficiency to 28 cents.

Proportional subsidy, insurers are allowed to adjust bids The idea of these counterfactuals is similar to the logic we discussed in the partial equilibrium framework. Proportional subsidies have two effects relative to the observed mechanism. First, there is a price level effect, by which, for example, a very generous subsidy would decrease the overall level of prices. Second, there is a significant change in relative prices that makes the more generous plans relatively more attractive.

We report the results for two proportional subsidy mechanisms in columns two and three of the upper panel in Table 2.5. We set the premium to be equal to 5% and 32% of the bids in these counterfactuals:

$$p_{jt} = x * b_{jt}. \tag{2.20}$$

The idea of the 5% counterfactual is to test how insurers respond if we almost entirely subsidize the market and remove most of the price sensitivity, as the government is not price sensitive in the model. This counterfactual has the expected effect according to which the insurers dramatically increase their bids. We estimate an increase in bids on he order of 400%. Consequently, even though the individuals now pay only 5% of the bids, the premiums are still relatively high - at 238 USD a year on average. This drop in premiums increases enrollment almost three-fold to 80% in the PDP plans from the observed 28% baseline. Consumer surplus increases from slightly more than 1 billion to almost 9 billion. The change in insurers' profits is, as expected, even more substantial with an almost ten-fold increase. Given the dramatic increase in bids, government spending also increases substantially. The result is a drop in welfare levels by 60 billion USD. At the same time, the return on a dollar increases three-fold to 56 cents. Essentially, in this counterfactual, we generate a large transfer from the taxpayers to the insurers, with a less than one-to-one pass-though to consumers.

The counterfactual where we set the premium to be 32% of the bid is interesting, as this closely corresponds to the current mechanism. Under the current mechanism, the premium is equal to 32% of the average bid plus any difference between the bid and the average. With setting up a direct proportional subsidy, we remove the distortion in relative prices. The result of this equilibrium is a 4 percentage point increase in enrollment relative to the observed baseline and an increase in consumer surplus and producer profit. Consumer surplus increases despite the increase in average premiums, as removing the distortion in relative prices improves the allocative efficiency. While total welfare decreases by 5 billion USD, we get almost twice as much surplus per dollar spent.

2.6 Conclusion

As private markets are increasingly introduced into social insurance systems, the government faces the challenge of regulatory design that would maximize the welfare of these programs subject to the incentive constraints of the private insurers. In this paper we study the regulatory mechanisms through which the government currently administers subsidies in a large prescription drug program for US seniors - Medicare Part D. We show that the current supply-side regulatory mechanisms that tie the premiums and subsidies to the realization of average "bids" by insurers, as well as utilize the same "bids" to determine which plans are eligible to enroll low-income beneficiaries, introduce

efficiency-decreasing distortions into the market.

Using the data from the first six years of the program, we estimate an econometric model of supply and demand that incorporates the regulatory pricing distortions in the insurers' objective functions. We then conduct counterfactual analyses of what the premiums, enrollment, and surplus would be on this market under different (simpler) ways of providing the subsidies to Medicare beneficiaries. Our findings suggest that relative to the existing outside options hat we keep unaltered in the counterfactuals, the current levels of subsidies are generating small return on government spending. The results of counterfactual equilibria in which we simulate subsidies of 5% to 95%, demonstrate that the higher is the subsidy, the higher is the incentive for the insurers to increase their bids in the program. While the goal of subsidy programs is to ensure the affordability of insurance, in practice in oligopolistic markets, the insurers with market power will have the incentives to inflate prices and pass them through to the price-insensitive government. In less extreme examples, we conclude that proportional subsidies that are not tied to the bids may generate higher consumer surplus and higher payoffs than average-tied subsidies that alter relative prices. At the same time, similar payoff on public funds and higher levels of welfare can be achieved with the optimally chosen flat subsidies that are specified ex ante to the annual clearance of market prices.

Figure 2-1: Parameters of the Minimum Coverage Requirements in Medicare Part D

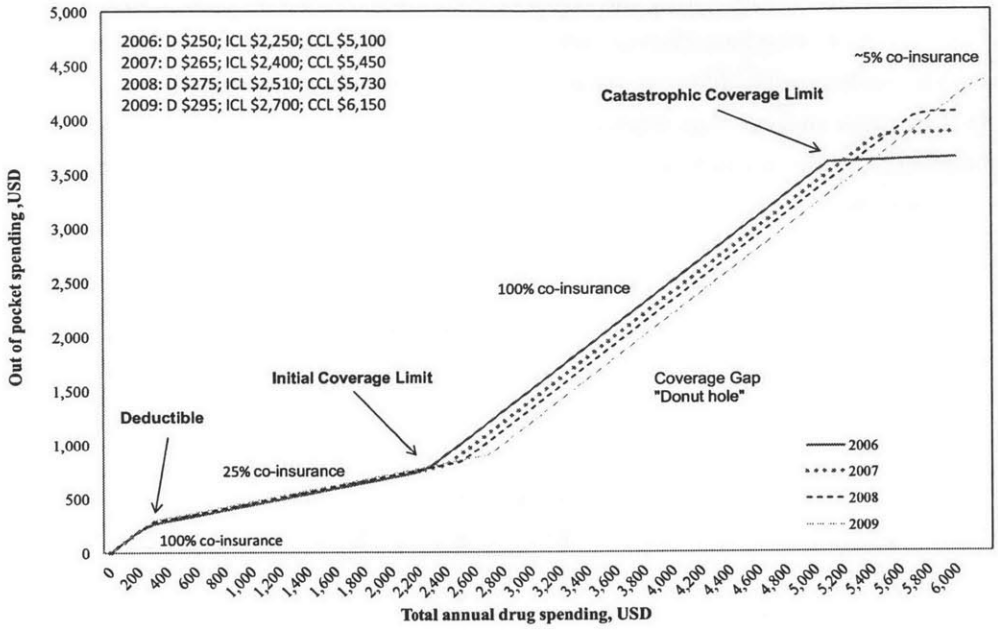

Insurers in the Medicare Part D program are required to provide coverage that gives at least the same actuarial value as the Standard Defined Benefit (SDB). The SDB design features a deductible, a co-insurance rate of 25% up to the initial coverage limit (ICL) and the subsequent "donut hole" that has a 100% co-insurance until the individual reaches the catastrophic coverage arm of the contract. The graph illustrates these features of the SDB by mapping the total annual drug spending into the out-of-pocket expenditure. Consider an individual, who in 2006 was in an SDB contract, and purchased prescription drugs for a total of $3,000. Out of this amount, the individual would pay the deductible of $250, then 25% of the next $2,000 up to the ICL of $2,500, and then 100% of the next $750 in the gap, for a total out of pocket spending of $1,500.

Figure 2-2: Cross-sectional variation in market shares of plans and outside option, 2010

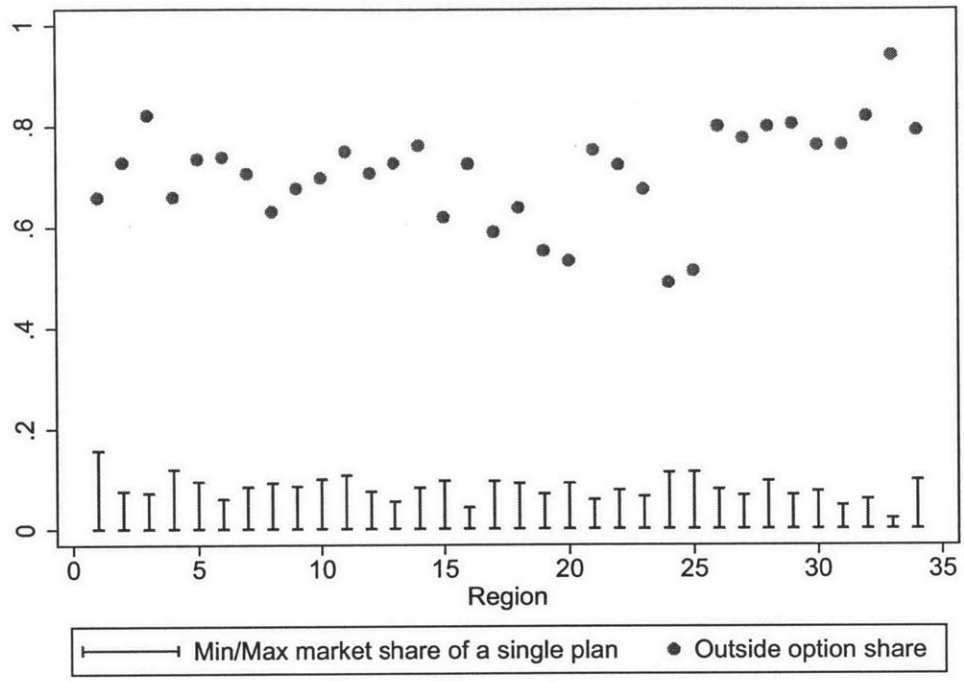

The figure depicts the distribution of market shares of the outside and inside options across 34 Medicare Part D regions. The definition of the outside option crucially affects our interpretation of welfare calculations. According to the Centers for Medicare and Medicaid Services, in 2010, almost 47 million individuals in the US were eligible for Medicare Part D coverage. This total number includes beneficiaries that are eligible both due to old age and because of disability (of which there are about 9 million), in all income groups, including individuals eligible for low-income subsidies. Out of this total number of eligible individuals, CMS estimates that about 42 million had what is called "creditable drug coverage" or a prescription drug coverage that satisfies a certain minimum requirement. Out of the 42 million with creditable coverage, 18 million were enrolled in stand-alone PDPs, 10 million were enrolled in MA-PDs, about 6 million received their Part D subsidies through their employer and about 8 million had coverage through other sources, such as federal employee or military insurance. In the current graph, we focus only on regular enrollees without low-income subsidies. We include both individuals that are eligible for Medicare because of old age and for other reasons. We consider all non-PDP sources of coverage to be in the outside option. The inside option market shares only include PDP plans.

Figure 2-3: Cross-sectional and time-series variation in observed contract premiums

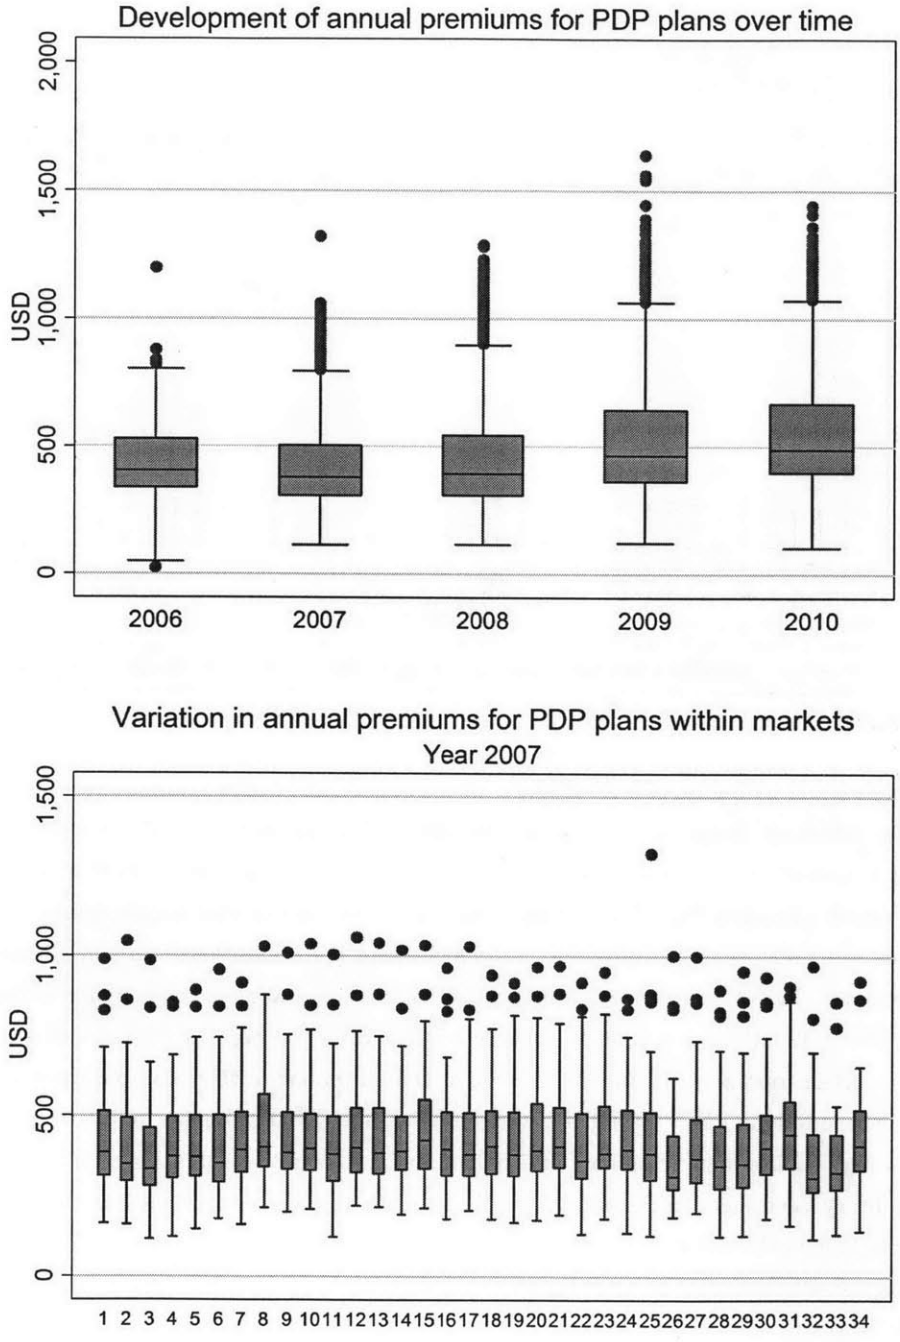

This figures illustrate the cross-sectional within and across market variation in premiums, as well as the time-series development of the cross-sectional variation. The premiums are scaled to annual dollar amounts. The data includes the stand-along PDP plans that we use in the estimation of the empirical model for years 2006-2010. Premiums are not weighted by enrollment.

Figure 2-4: Premiums and market structure

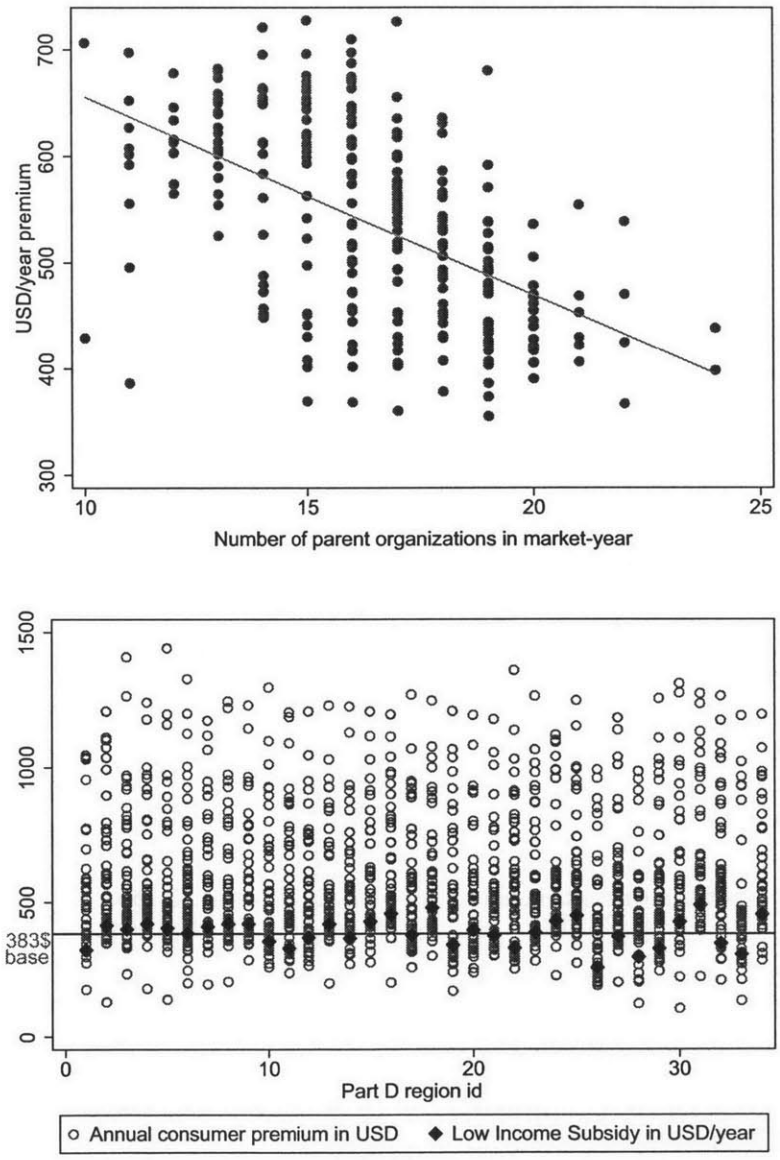

Panel 1. Data shows a negative correlation between the number of competing insurers and average premiums in the market. Panel 2. Since Medicare uses both MA-PD and PDP plans to administer subsidy calculations, we observe that the base premium is low relative to the distribution of PDP premiums. Further, only a small fraction of plans qualify to enroll beneficiaries with low-income subsidies.

Figure 2-5: Marginal cost projection and prediction for manipulating and non-manipulating plans

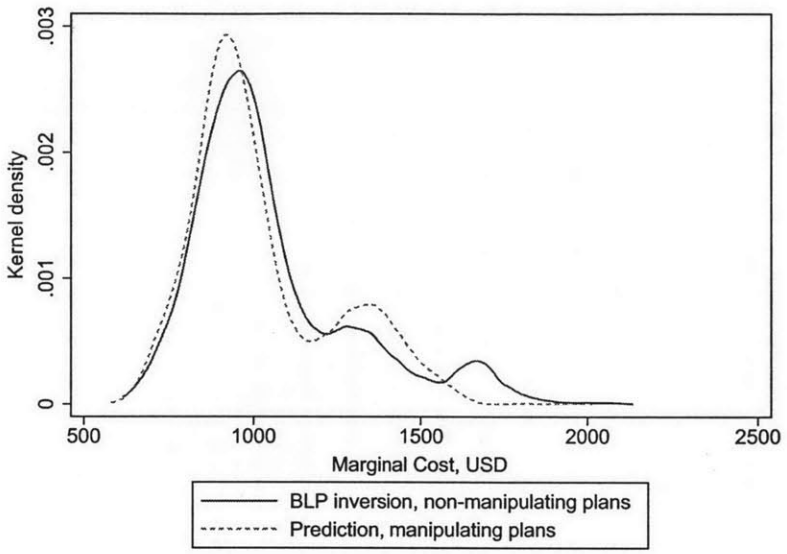

	Plan group 1	Plan group 2
Plan characteristics	Estimated MC	Estimated MC
Deductible	-0.00623**	-0.00431***
	(0.00177)	(0.000696)
Coverage in gap	3.770***	4.137***
	(0.307)	(0.422)
Basic plan	0.812*	0.517**
	(0.272)	(0.112)
Share of top drugs covered	-6.040	0.766
	(3.782)	(2.781)
Share of unrestricted drugs	0.00400*	-0.000224
	(0.00168)	(0.00159)
Number of preferred pharmacies	0.000017***	0.000012
	(1.90e-06)	(5.33e-06)
Unobservable component	0.438***	0.424***
	(0.0635)	(0.0614)
Constant	11.53**	10.99***
	(2.369)	(1.711)
Observations	818	2,344

Figure 2-6: Counterfactual Consumer Welfare as a Function of Subsidy Proportion

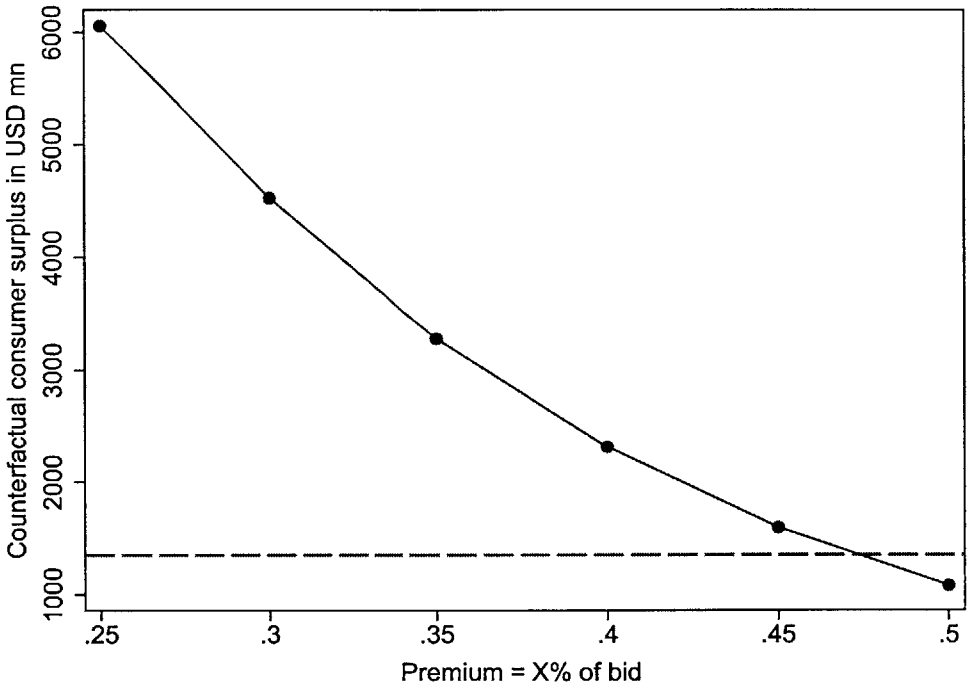

The figure illustrates how the estimated consumer surplus changes if we keep the bids submitted by Part D plans fixed, but vary the percentage of bid paid by the beneficiaries in premiums. Note that in this partial equilibrium framework, a direct proportional subsidy of about 68% (implying the individual premium that is equal to 32%) would result in much higher consumer surplus than the current subsidy mechanism, by which the premium is equal to 32% of the average bid plus the difference between the bid and the average. The estimated consumer surplus at current premium levels and under the observed allocation is indicated by the horizontal line. The intersection indicates the level of proportional subsidy that would achieve the same level of consumer welfare. These results suggest that the current subsidy mechanism is distorting the relative attractiveness of Medicare Part D plans.

Figure 2-7: Partial Equilibrium: Premium = Full Bid

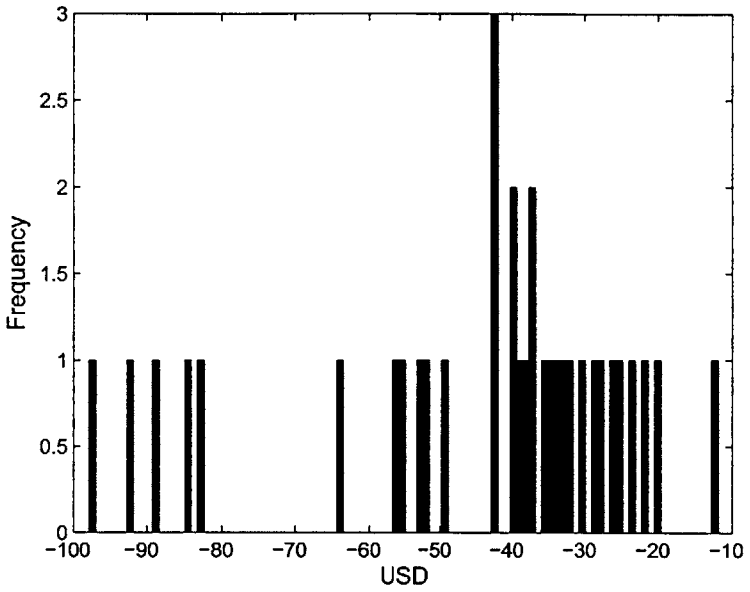

Figure 2-8: Partial Equilibrium: Premium = Bid-min(Bid)

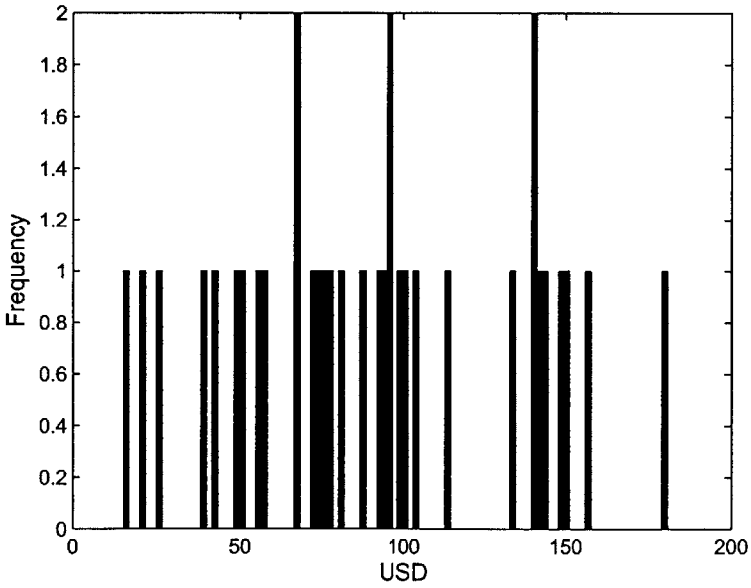

The figure shows a counterfactual distribution of average change in consumer surplus across 34 Medicare Part D regions in 2010. The change in the counterfactual consumer surplus is calculated relative to the surplus estimation for the allocation at prices observed in the data.

Figure 2-9: Partial Equilibrium: Premium = 25 Percent of Bid

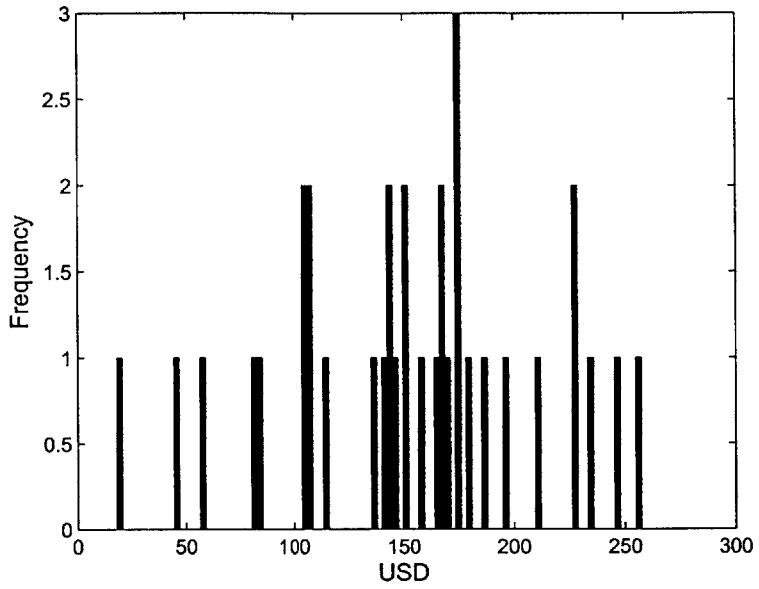

Figure 2-10: Partial Equilibrium: Premium = Marginal Cost

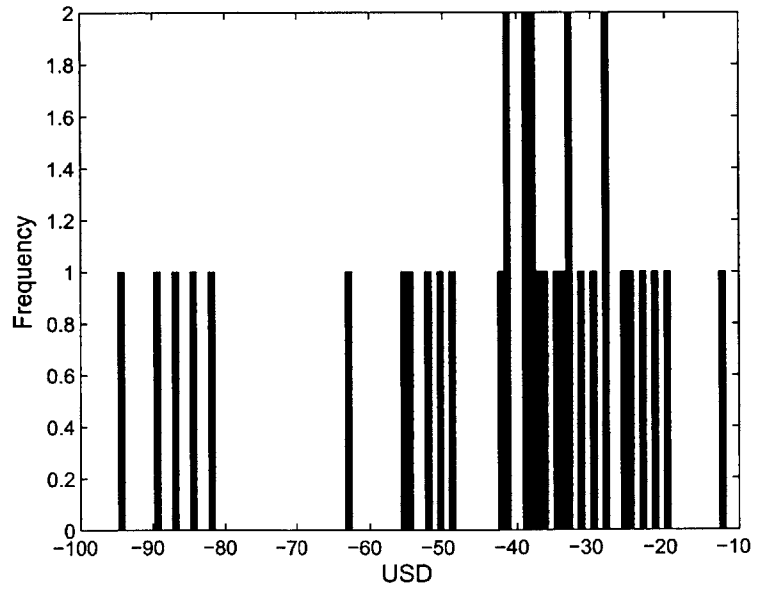

The figure shows a counterfactual distribution of average change in consumer surplus across 34 Medicare Part D regions in 2010. The change in the counterfactual consumer surplus is calculated relative to the surplus estimation for the allocation at prices observed in the data.

Table 2.1: Summary Statistics for the Estimation Sample of PDP Plans

Year	2007	2008	2009	2010
Plans				
Total number of plans	1,574	1,590	1,462	1,542
Average number of plans per market	46	47	43	45
Firms				
Total number of parent organizations	54	52	50	49
Average number of parent organizations per market	19	19	18	17
Premiums				
Unweighted average annual consumer premium	$423	$475	$536	$562
Unweighted average premium per market	$419	$471	$526	$552
Subsidies				
CMS national average bid (annual)	$965	$966	$1,012	$1,060
CMS base consumer premium (annual)	$328	$335	$364	$383
CMS subsidy for average risk beneficiary	$637	$631	$648	$677

Table 2.2: Demand Estimates

	OLS	2SLS	BLP	BLP MicroMoments
	(1)	(2)	(3)	(4)
Annual premium, in 100 USD	-0.368***	-0.654***	-1.931***	-1.899***
	(0.02)	(0.07)	(0.08)	(0.08)
Annual premium, Sigma	-	-	0.646***	0.634***
			(0.08)	(0.10)
Annual premium x Income	-	-	0.001	0.002
			(0.01)	(0.01)
Annual premium x De-meaned Age	-	-	-0.019	-0.012
			(0.21)	(0.22)
Deductible, in 1000 USD	-3.523***	-4.868***	-6.935***	-6.890***
	(0.25)	(0.42)	(0.50)	(0.49)
Coverage in the gap	55.09***	163.0***	260.244***	269.852***
	(6.58)	(25.56)	(30.63)	(29.96)
Coverage in the gap, sigma	-	-	0.756	-2.48
			(486.12)	(473.49)
Coverage in the gap x Income	-	-	-0.236	-0.67
			(11.48)	(11.54)
Coverage in the gap x De-meaned Age	-	-	-0.084	1.93
			(112.21)	(119.96)
Number of most common APIs covered	1.381**	2.236***	9.754***	9.6702***
	(0.62)	(0.72)	(1.45)	(1.47)
Number of in-network pharmacies	0.715***	0.643***	0.432***	0.449***
	(0.27)	(0.23)	(0.12)	(0.12)
Plan vintage	66.74***	77.89***	95.893***	95.447***
	(4.14)	(4.56)	(3.78)	(3.80)
Insurer FE	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Observations	6,024	6,024	6,024	6,024
R-squared	0.478	0.443	-	-

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table shows four sets of demand estimates. Each estimation uses data on Medicare Part D stand-alone prescription drug plans in years 2007 to 2010. In addition to the displayed coefficients and fixed effects, all regressions also include a constant and the following plan characteristics: a dummy for an enhanced plan; number of APIs in formulary; number of drugs placed in Tiers 1-2 of the formulary (i.e. having low cost-sharing). 2SLS specification of the Berry Logit uses a collection of traditional BLP-style and Hausman-style instruments, see the text for more details. The BLP specifications with and without micro-moments in addition use Hausman instruments interacted wight the mean and standard deviation of the income and age distributions in each region/year. The last specification includes micro-moments generated by constructing the share of individuals above and below the age of 75 buying a plan with any coverage in the gap in each region-year combination. Robust standard errors are brackets. *** p<0.01, ** p<0.05, * p<0.1

Table 2.3: First stage and Reduced Form - BLP and Hausman instruments for premiums

	First Stage	Reduced Form
	Monthly premium	Log shares
Hausman instrument	0.0449***	-0.00403***
	[0.00339]	[0.000307]
No. similar plans by PO in the same region	-0.951***	0.128***
	[0.191]	[0.0245]
No. of MA plans in a region-year	-0.00563	-0.00119
	[0.0151]	[0.00100]
No. of MA plans in a region-year by same PO	-0.0339*	0.00913***
	[0.0198]	[0.00246]
No. of common APIs covered by MA plans in region	0.000853**	0.000352***
	[0.000371]	[6.20e-05]
Average deductible of MA plans in the same region	-0.0141	0.00261**
	[0.0143]	[0.00125]
No. of top drugs by MA plans in the same region	0.0341***	-0.00180***
	[0.00672]	[0.000585]
No. of common APIs covered by other PO in region-year	0.00377***	-0.00113***
	[0.000568]	[0.000107]
Observations	6,024	6,024
R-squared	0.77	0.526
F(8,135)	30.8	

The table reports the first stage and the reduced form for variables that are used as instruments for premiums. Each regression uses data on Medicare Part D stand-alone prescription drug plans in years 2007 to 2010. See the text for more details on the construction of Hausman and BLP-style instruments. Robust standard errors are in brackets. PO stands for Parent Organization. MA stands for Medicare Advantage. API stands for Active Pharmaceutical Ingredient. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 2.4: Counterfactual Welfare Estimates: Partial Equilibrium and Marginal Cost Pricing

	Current Mechanism	Partial Equilibrium Counterfactuals			MC Pricing
	P=data	P = bid	P = bid-min(bid)	P = 25% of bid	P=MC
Surplus calculations (USD billions)					
Consumer surplus	1.35	0.02	4.11	6.05	0.05
Producer profit pre-reconcilliation and selection adjustments	0.87	0.01	1.56	3.02	-
Government transfers on premium subsidies	5.94	-	15.48	25.01	-
Government transfers on subsidies other than premiums	3.44	0.05	7.25	11.04	0.11
Total welfare, including cost of public funds	(9.97)	(0.04)	(23.88)	(37.78)	(0.10)
Value generated per 1 USD of government funds	**0.18**	**0.45**	**0.19**	**0.19**	**0.32**
Enrollment and prices					
Enrollment of regular beneficiairies (millions)	8.77	0.13	17.07	21.58	0.29
Inside option enrollment share	28%	0.4%	55%	69%	0.9%
Enrollment-weighted average annual premium (USD)	502	1,183	275	386	1,079

The table reports counterfactual simulations of different subsidy mechanisms in the partial equilibrium setting. This implies that we keep the submitted bids of insurers as observed in the data and only change the mechanism through which these bids are transformed into beneficiary premiums. It is important to emphasize that this table represents the *lower* bound of values generated per dollar of government spending, as in these calculations we are assuming that the government does not have expenditures if individuals leave PDP plans to the outside option. In practice, if individuals substitute to other government programs, the spending reduction (or in other words, the opportunity cost of funds) may be significantly lower. In future work, we will develop the upper bounds of these measures.

Table 2.5: Counterfactual Welfare Estimates: Full Equilibrium

	Current Mechanism	Full Equilibrium Counterfactuals		
	P=data	P=5% of bid	P=32% of bid	P=95% of bid
Surplus calculations (USD billions)				
Consumer surplus	1.35	8.68	2.15	0.02
Producer profit pre-reconcilliation and selection adjustments	0.87	84.50	5.01	0.02
Government transfers on premium subsidies	5.94	112.45	11.99	0.01
Government transfers on subsidies other than premiums	3.44	14.80	5.28	0.06
Total welfare, including cost of public funds	(9.97)	(72.25)	(15.30)	(0.04)
Value generated per 1 USD of government funds	**0.18**	**0.56**	**0.32**	**0.56**
Enrollment and prices				
Enrollment of regular beneficiairies (millions)	8.77	24.83	9.97	0.13
Inside option enrollment share	28%	79.7%	32.0%	0.4%
Enrollment-weighted average annual premium (USD)	502	238	566	1,175

	Current Mechanism	Full Equilibrium Counterfactuals		
	P=data	P=bid-0.68(av.b)	P= bid - $676	P=bid-$1340
Surplus calculations (USD billions)				
Consumer surplus	1.35	2.70	1.18	11.03
Producer profit pre-reconcilliation and selection adjustments	0.87	1.97	0.95	6.86
Government transfers on premium subsidies	5.94	10.13	4.06	36.55
Government transfers on subsidies other than premiums	3.44	5.28	2.60	12.80
Total welfare, including cost of public funds	(9.97)	(15.36)	(6.51)	(48.27)
Value generated per 1 USD of government funds	**0.18**	**0.23**	**0.25**	**0.28**
Enrollment and prices				
Enrollment of regular beneficiairies (millions)	8.77	12.11	6.00	27.28
Inside option enrollment share	28%	38.9%	19.3%	87.6%
Enrollment-weighted average annual premium (USD)	502	398	555	78

The table reports counterfactual simulations of different subsidy mechanisms in the full equilibrium setting. It is important to emphasize that this table represents the *lower* bound of values generated per dollar of government spending, as in these calculations we are assuming that the government does not have expenditures if individuals leave PDP plans to the outside option. In practice, if individuals substitute to other government programs, the spending reduction (or in other words, the opportunity cost of funds) may be significantly lower. In future work, we will develop the upper bounds of these measures.

Chapter 3

Cream-Skimming and Enrollment Mandate in Public-Private Health Insurance: Evidence from Germany

3.1 Introduction

The ubiquitous feature of health insurance markets is that the insurers' costs depend on who their enrollees are and how they behave. This feature of selection markets raises concerns about the feasibility of efficiency-improving competition in health insurance, and has served as a traditional rational for extensive government intervention in health insurance. Increasingly, public policies in healthcare attempt to strike a balance between the selection concerns and the efficiency advantages of competitive markets by reorganizing the purely public or purely private health insurance systems into different mixtures of the two. A long-standing question in such arrangements, where a private system exists in parallel to a public one, is whether the private insurers successfully cream-skim good risks outside of the public option. The second question is whether there are policies that can effectively address this concern.

In this paper, I investigate these issues in a unique institutional set-up of the German health insurance market that is characterized by the co-existence of two competing insurance systems. The two systems differ in the ability of insurers to underwrite individual-specific risk and in the ability to offer contracts with different cost-sharing structures. While the non-profit insurers in the so-called "statutory" insurance system have to offer the federally set income-adjusted community rating premiums, and face a regulatory lower-bound on the generosity of coverage, the competing private insurance sector is less restricted in coverage levels, and its prices are regulated only to the

[0]I am indebted to Amy Finkelstein and Stephen Ryan for their guidance throughout this project. I also thank the participants at the MIT Public Finance and Industrial Organization lunches, 15th IZA European Summer School in Labor Economics, and MEA Seminar at the Max Planck Institute for Social Law and Social Policy for their feedback. Data for this project - the Scientific Use Files of the German Socio-Economic Panel (1984-2007) - were provided by DIW Berlin and Cornell Department of Policy Analysis and Management, which I gratefully acknowledge.

extent that they have to follow annuity-like long-term contracts. Consequently, while individuals with the same income pay the same premiums in the statutory system independently of their risk, enrollees of the private plans face full underwriting of their individual risk and may be rejected by the insurers.

Using several unique features of this institutional environment, the paper accomplishes several tasks. First, I test for the presence of cream-skimming of good risks out of the public system, using a fuzzy regression discontinuity design to separate selection from moral hazard. To begin with, I document several patterns in the data that suggest little difference between the observed healthcare utilization or health risk (based on diagnostic patterns) between individuals that self-select to opt out of the statutory system and enroll with a private insurer. This evidence tests for the presence of asymmetric information that may combine selection and moral hazard effects. I then use an RD design to isolate the causal effect of insurance on healthcare utilization, or moral hazard. This RD design is based on an income threshold that regulates which individuals are required to be enrolled with the SHI and which are allowed to choose between the statutory and the private system. Specifically, employees, whose income is below about 50,000 EUR/year are mandated to enroll in the statutory system and cannot switch to one of about 40 independent private insurance providers. The RD-based estimates of moral hazard suggest no economically significant causal effect of having private insurance on healthcare utilization. The estimate of moral hazard in conjunction with the estimates of differences in healthcare utilization across the self-selected samples, allow me to calculate residual asymmetric information coming from selection. Overall, I find no evidence that private insurers manage to cherry-pick good risks out of the public system. I discuss two possible explanations for this result.

First, as has been argued in the literature that analyzed selection in other insurance markets, the presence of heterogeneous preferences on dimensions other than expected spending, may imply that there is no clear relationship between the insurance contract chosen and the individual's risk type (Finkelstein and McGarry, 2006). I take advantage of the detailed survey data and estimate individual preferences for different types of insurance. I find indeed that individual characteristics plausibly not directly related to expected health spending, such as for example political views, play an important role in the observed insurance choices. The second hypothesis concerns the supply-side of the market and argues that the lack of adverse selection may be driven by the inability of private insurers to cream-skim effectively because on the long-term annuity structure of their contracts interacted that incentivizes the individuals to enroll into the system as early as possible in their lifetime, at which point the insurer has only very limited scope for underwriting. The finding of no selection on the extensive margin between the statutory and the private insurance systems suggests that removing the statutory enrollment mandate would increase consumer surplus *ceteris paribus*. Using the estimated preferences for private insurance, I am able to quantify this change in consumer surplus.

The paper is related to several strands of literature. First, it is related to the extensive literature that tests for the existence of adverse selection in insurance markets. Einav, Finkelstein, and Levin

(2010a) provide a detailed survey of this literature. One observation in this literature is that we sometimes fail to find evidence of adverse selection even in environments where insurance contracts clearly differ in their generosity of coverage. A possible explanation for these findings is the existence of heterogeneous preferences for health insurance coverage above and beyond the heterogeneity in risk types. This observation has spurred a strand of literature that studies the role of heterogeneous preferences in selection markets. The issue of heterogeneous preferences in the demand for health insurance is discussed in Cutler, Finkelstein, and McGarry (2008). They point out that in practice, low risk individuals do not necessarily buy less insurance coverage. They argue that individuals' preferences, such as for example the degree of risk aversion, may reverse the relationship between the risk type and the selected level of coverage. Geruso (2013)empirically explores the theme of heterogeneous preferences and finds that older individuals enroll in more comprehensive plans than younger individuals with the same expected healthcare expenditure risk.

Relatedly, the paper complements a relatively small, but growing literature on the nature of competition, contract structure, and the role of government regulation in competitive selection markets. This literature recognizes that the presence of adverse selection in an insurance market is a function of insurers' behavior as much as it is a function of individual choices. Insurers select which menu of contracts they offer to potential enrollees, and these choices determine the extent to which insurers are able to cream-skim profitable risks in the market. The insurers' choices of contract menus are in turn affected by their expectations about individuals' demand for different contracts, including all dimensions of preference heterogeneity and the underlying risk type, as well as by a multitude of government interventions that may alter the demand response or the space of possible contracts. In this vein, Lustig (2011) studies the interaction of adverse selection and imperfect competition on the Medicare+C HMOs market, allowing the insurers to endogenously select the generosity of their contracts. Kuziemko, Meckel, and Rossin-Slater (2013) explore whether health insurance companies compete on risk in the context of Medicaid managed care market. They find evidence consistent with a model of cream-skimming, where insurers try to retain low-cost enrollees and pass high-cost risks to competitor. Bauhoff (2012) conducts an audit study in Germany, documenting how insurers within the SHI system screen risks. The audit study design allows the paper to disentangle demand-driven self-selection from cream-skimming. Brown, Duggan, Kuziemko, and Woolston (2012) explore a setting that is close in spirit to the institutional framework considered in the current paper - they explore the selection of risks between the Medicare fee-for-service and the Medicare Advantage program. The paper finds that private insurers that participate in the Medicare Advantage program select risks even in presence of risk adjustment mechanisms.

Finally, this paper is closely related to the literature that has studied the German health insurance system. Nuscheler and Knaus (2005) address the issue of risk selection among different sickness funds that offer health insurance within the German statutory system, finding that the observed differences in the risk pools of different sickness funds are mostly due to the consumers' switching costs and that the sickness funds themselves don't practice any significant cherry-picking

of consumers. Using an RD design similar to the one I exploit in the current paper, Hullegie and Klein (2010) estimate that holding a private insurance policy decreases the number of doctoral visits, doesn't affect the number of hospital stays and improves self-assessed health. Schmitz (2011) uses a self-reported measure of risk aversion to show some evidence of advantageous selection into the supplementary coverage market. Grunow and Nuscheler (2013) study the issue of selection patterns between the private and statutory systems in Germany. Their paper argues that the private insurers are unable to select good risks at the enrollment stage, but manage to return high-risk individuals back to the public system later.

The rest of the paper is structured as follows. Section 2 discusses a stylized framework of the German health insurance market and the expected nature of selection and its interaction with regulatory interventions. Section 3 describes the data used in this study and provides empirical evidence about the nature of selection. Section 4 analyzes the demand for private health insurance, including the examination of the private insurers' pricing and welfare implications of the access restrictions given the nature of selection between the systems and demand for the PHI. Section 5 briefly concludes.

3.2 Conceptual framework

3.2.1 Stylized model of German health insurance system

Consider the following stylized model of the German two-tier health insurance market.

Suppose we have a set of beneficiaries that can be characterized by type θ_i. The individual's type is a vector of characteristics that includes expected healthcare utilization costs, risk preferences and, importantly, other potential sources of heterogeneous preferences for different insurance systems. The latter may include preferences for premium services or convenience in healthcare consumption, such as preferences for larger physician networks or shorter waiting times. Assume that every individual is of a different type and that there is some distribution of these types in the population $\theta_i \sim F(\theta)$. The assumption of a multitude of heterogeneous preferences is in line with the recent literature that has identified and measured the effects of preferences on health insurance choices. For example, (Geruso, 2013) finds residual health insurance preferences related to age. The literature has also found preference heterogeneity in other insurance markets, for example in annuities (Einav, Finkelstein, and Schrimpf, 2010b) and long-term-care insurance (Finkelstein and McGarry, 2006) . These studies predominantly analyze how different preferences affect the extent of coverage that individuals choose to buy. In other words, the analyzed insurance contracts lie on a vertically differentiated product spectrum, and individuals agree on the ordering of contract valuation, but not on its level. A preference for more *convenience* rather than coverage in health insurance that is introduced in this paper, has not, to the best of my knowledge, received much explicit attention in the previous literature.[1] Introducing this preference allows for a horizontal differentiation of

[1]The closest to the concept of the horizontally differentiated plans is the comparison between the HMO and PPO policies that has been studied in the literature.

insurance contracts, since the degree of coverage is interacted with differential convenience, which may be valued differently by individuals.

Given the vector of consumer types, both individual expected medical expenditure risk and preferences can be parsimoniously summarized in a static utility framework. Specifically, suppose that individual's utility from any of the insurance plans offered on the market takes the following form[2]:

$$U_{ij} = u_{ij}(\theta_i) - p_{ij}(\theta_i)$$

Here, individual i gets some value $u_{ij}(\theta_i)$ from enrolling into an insurance plan j. This valuation depends on the individual's type θ_i and the type of the chosen insurance plan. The difference in the valuation between the the plans comes from heterogeneous preferences for the level of services, benefits, and cost-sharing. The individual has to pay price p_{ij} for the chosen insurance plan. Unlike in many standard consumer goods markets, the price for the insurance contract is allowed to depend on the individual's type. The individual chooses insurance product j, so as to maximize utility.[3]

Now consider the supply side of the market. The stylized description of the health insurance market in our setting involves two firms P [PHI] and S [SHI].[4] The firms sell health insurance products that are horizontally differentiated. While the common perception is that the private insurer provides higher quality products, at the same time it also offers different cost-sharing arrangements.[5] Therefore, the consumer ranking of the products may disagree depending on individual

[2]We may instead start with a dynamic life-cycle model of insurance choices, where individuals choose their preferred insurance coverage in every period given heterogeneity in preferences and uncertainty about the expected healthcare expenditure. However, given the institutional nature of choice, I view it as appropriate to approximate the dynamic model with a static framework. Under the current regulation, individuals basically have one-in-a-lifetime decision (if any) of which insurance plan to enroll into, since individuals that opt into the PHI system are restricted in the ability to switch back to the SHI. In that case, the expected continuation value in the Bellman equation of choosing the PHI folds into a constant. Further, for the individuals in SHI that eventually want to switch to the PHI, any waiting is suboptimal, given the annuity nature of the PHI. Thus, a static framework, where the valuations refer to life-time utility from being in one or the other system, seems to be an appropriate approximation to the choice problem at hand.

[3]Here I am abstracting from a number of institutional details of this choice. Thus, for instance, the concept of price for the contract in the static framework is not simple. First, in reality individuals pay monthly insurance premiums and these premiums change annually both on the SHI and the PHI contracts. The SHI prices usually rise because of the changes in regulation and any changes in income. The PHI prices may change due to higher costs of service, demographic, or regulatory changes. Further, PHI prices depend on the individual's entry age, since the contract is structured as an annuity. The static framework provides the closest approximation if the prices $P_{ij}(\theta_i)$ are viewed as life-time expected prices the individual faces for his or her insurance choices. Since the regulatory constraints are such that the choice of insurance can be approximated as one-in-a-lifetime decision, the interpretation of the valuations and prices as life-time variables appears appropriate. In the empirical work, however, I will have to resort to monthly premiums and valuations, since the data provides only short snapshots of the individual's experiences with their insurance choices.

[4]A natural concern here is of course that both the PHI and the SHI systems actually consist of many separate firms. Since I do not have data on the choices of specific PHI providers, these choices cannot be accounted for empirically and therefore are also abstracted from in the conceptual discussion. Therefore, one could think of the presented choices as being at the top level of a nested choice problem.

[5]Even for PHI contracts with zero deductible and zero co-insurance, the PHI feature of reimbursing individuals

preferences. Assume that both firms have no administrative cost, so that their only cost are the expenses caused by consumers' healthcare utilization. In this setting, the cost that each firm experiences for covering consumer i with characteristics θ_i depends on the consumer's characteristics and the quality of contract that the consumer chose. Let us denote the incurred expenditures for a given healthcare service and a given contract quality with $e_{Pi}(\theta_i)$ for an individual with policy P and $e_{Si}(\theta_i)$ for an individual with policy S. Since firm P reimburses more to providers, its initial expenditures for the same consumer and for the same service are always higher than those of firm S. That is, $e_{Pi}(\theta_i) > e_{Si}(\theta_i)\ \forall i$.

However, the actual cost for firms P and S to cover an individual i in a given time period, depend on the cost-sharing arrangements of the individual's contract. The cost-sharing arrangements both decrease the cost for the firms by directly passing a part of the cost to the consumer and may decrease the cost indirectly if the individual consumes less services due to cost-sharing. Since firm P tends to have both more cost-sharing and higher provider reimbursement, the net difference between the cost of coverage for a given individual i between the two firms is ambiguous. That is,

$$c_{Pi}(\theta_i) \gtreqless c_{Si}(\theta_i)\ \forall i$$

Although the comparison of costs between the two insurance providers is ambiguous, each firm knows its terms of contract with individuals and thus has some information about its expected costs on which it could base its pricing strategies. These pricing strategies are constrained by the government's regulation. The prices that both firm offer have to depend on their expected cost of coverage. This cost depends on the quality of coverage that the firm provides. To emphasize this relationship, I include the quality variables q_s and q_p into the cost functions.

The quality of firm S is strongly regulated and thus firm S does not make choices of quality. Furthermore, I suppose that firm S has to set its price to the average costs that it faces given its customer base weighted by income profiles in the population ω_i. That is, the price that individual i faces for policy S can be expressed as:

$$p_{Si} = \omega_i \int\limits_{\{\theta|buy\,S\}} c_S(\theta, q_s) d\theta$$

I assume that firm P gets to set the price after it observed the regulated price and quality of firm S. That is, firm P can respond strategically to the price set by firm S. In particular, if firm P competes with firm S on selection, then it would set its prices so as to attract the most profitable consumers. Further, firm P can price-discriminate on the observable part of consumers' characteristics, which makes the strategic response to the price of firm S easier to execute. Since the price that firm P offers to consumer i has to depend on expected costs, given the observable characteristics of the

with 1 or 2 monthly premium payments for no-claims years, encourages the PHI-insured to not submit small claims and thus have a de facto deductible.

consumer, the price will be a function of expected costs and the price offered by firm S:[6]

$$p_{Pi} = f(c_{Pi}(\theta_i^{observable}, q_P), p_S)$$

In the equilibrium of this game, under the assumption that firm S sets its quality to the minimal regulatory constraint, the key choice variable is the quality provided by firm P. This quality (which refers to the combination of benefits and cost-sharing) determines the relative prices and sorting of consumers across the two systems. The next section discusses the possible implications of different pricing strategies in the two systems on the directions of risk sorting.

3.2.2 Cream-skimming and enrollment mandates

Given the differences in pricing methods, different types of individuals may be considered "good" or "bad" risks for PHI vs. SHI and we would expect these differences to impact selection patterns. Specifically, for the employees above the income threshold, the SHI charges a fixed premium. Thus, the "good risks" for the SHI are simply those individuals whose healthcare utilization expenditures in this year are lower than what they pay into the system. Let us call these individuals "net payers" and the individuals that spend more on their healthcare than they pay, "net receivers."[7] Then, I can define selection in this market for the purpose of this paper as follows. There is adverse selection into the PHI if the individuals that opt out for the PHI would have been predominantly "net payers" in the SHI system. There is advantageous selection into the PHI if the individuals that opt out for the PHI would have been "net receivers" in the SHI system. And finally there is no selection if the switchers are a random mix of risks. Especially in the case of the adverse selection, we would be tempted to claim that the PHI competes with the SHI purely on selection rather than on efficiency or better quality. In this case, competition may appear to be harmful; while in the case of advantageous selection, competition would appear to be welfare-improving.

Before looking for any evidence of selection going one way or the other in the data, it is useful to explore whether we would intuitively predict any natural direction for selection. For this, it is essential to closer understand the pricing scheme of the PHI and which kinds of individuals it would have an incentive to select. The PHI premium is calculated as an annuity on the expected life-time healthcare costs. The theoretical idea is to have fixed life-long monthly payments for the individual, while the insurance company absorbs any risk on the variation in the healthcare costs above the predicted levels. Because of the annuity structure, the monthly premium payments are usually higher than the claims in the younger age and the difference between them is held (invested) by the insurance company and then used to pay for the higher expenses in the old age. That is, the very stylized calculation of the premium is to take the PDV of the expected lifetime healthcare spending and divide them equally across all expected life months.

[6] I discuss the actual pricing technique of the private insurance providers in more detail in section 5.1

[7] An alternative definition would be to consider life-time payments into the system versus lifetime healthcare expenditures and classify the individuals according to this criterion. Given that the SHI providers need to balance their annual budgets, the annual classification of risks appear more natural.

This calculation faces two problems that do not allow the insurance companies to actually maintain the monthly premium fixed over the individual's lifetime. First, the society is aging and so the assumptions about life expectancy at the entry age will be underestimating the number of years the insurance will likely have to cover. This is especially critical, since the additional years of life in the old age are probably going to be the most expensive years in terms of the medical spending. Second, the prediction of the expected claims in the older age is based on the current expenditures of the company's customers at this age. That is, to predict how much spending a current 30-year male will have when he is 50, the company looks at the realized average spending of current 50-year old males in the same or similar insurance plan. Given the fairly steep growth in the medical expenditures, this premium calculation severely underestimates the actual expected costs for our 30-year old. To cope with these two issues, the insurer has to annually adjust up the premiums for the existing contracts.[8]

With this pricing mechanism in place, let us consider which individuals constitute "good risks" for the PHI and how these compare to the "net payer" and the "net receiver" classification in the SHI as suggested above. The PHI uses average costs by age and gender in a given plan as the basis for the calculation of the premiums for a new consumer. The insurer can also add individual-specific risk adjustment payments or exclude pre-existing conditions from coverage. A "good risk" for a specific PHI plan is then someone, whose healthcare expenses over lifetime are lower than the average expenditures that the insurer uses to calculate the premium. It follows that in the pool of employees above the income eligibility threshold in a given year, the PHI would try to select individuals that are expected to have healthcare utilization below the average utilization in that PHI plan.

At the same time, however, given the annuity system of the PHI, younger individuals should find it relatively more attractive to join the PHI system than older individuals. Thus, if PHI observes all its applicants at a relatively young age,[9] it may be hard for the insurer to predict how the healthcare utilization of a specific individual will relate to the average in the plan. That is, in this setting the level of informational uncertainty is so high that it may prevent the insurer from successfully executing any meaningful selection of good risks. In other words, the very nature of the PHI pricing scheme and the additional regulation that essentially prohibits the PHI to terminate existing contracts, dampens the extent of targeted selection.

Furthermore, even if PHI could select and if we viewed both systems as charging fixed monthly premiums over the life-time of the individual (and this would have been true absent the growth in healthcare costs that in principle affects both systems), the classification of "good" and "bad"

[8]This discussion of the pricing methodology is provided to the best of my knowledge and is based on external information about the operations of the PHI providers and the regulatory provisions for the industry; actual pricing strategies of individual insurance companies may be different. Specifically, the discussion provided here relies heavily on Kalkulationsverordnung KalV version 2009 and Schneider(2002) presentation at the 27th International Congress of Actuaries

[9]Furthermore, note that for the employees, the PHI applicants are at the same time going to be high-earners, implying that these individuals are less likely to have any significant chronic conditions or disabilities

risks across the systems would be different due to the different base used for premium calculations. That is, someone, who is a "net payer" in the SHI system, may be a "bad risk" for the PHI system, if this individual's healthcare costs are higher than the average for his/her age and gender in the PHI. Similarly, someone, who is a "net receiver" in the SHI, may still be profitable for an expensive tariff in the PHI. Thus, the nature of selection between the two competing systems as defined in the beginning of this section appears to be ambiguous a priori.

Understanding if and what kind of selection occurs at the intersection of the systems is policy-relevant in light of the regulatory restrictions on the access to the PHI through mandated SHI enrollment. Suppose that despite the theoretical ambiguity, we were to find empirically that there is adverse selection from the SHI to the PHI. Then, the access restriction policy in place would be ensuring that no Akerlof (1970) style unraveling can occur in the SHI structure, since the majority of individuals under this insurance coverage are in the non-selected risk pool.[10] The welfare-improving case of the access restriction would be significantly weaker, if advantageous selection were occurring. Lastly, consider the case of no selection. Suppose we believed that the reason for selection not occurring were indeed the pricing nature of the PHI, which induces individuals to switch to the PHI at as young age as possible. In that case, removing high-income eligibility threshold would allow individuals to apply for the PHI when they are even younger (assuming that most employees need some time to get to the high-earner status). This would make selection even harder for the PHI insurers, since they will face even less information about the applicant's expected risks. In this scenario, the current presence of the access restriction would again appear *ceteris paribus* welfare-decreasing.[11] To shed some light on the issue, the next section considers whether healthcare utilization data reveals any distinct selection patterns between the insurance systems.

3.3 Empirical evidence on the nature of asymmetric information

3.3.1 Data and descriptive statistics

Throughout the empirical analysis, I use data from years 2005-2009 of the German household survey panel SOEP. The survey offers a collection of self-reported answers for a sizable representative sample of the German population.[12] In the 2008 cross-section, the survey recorded observations on 18,703 individuals. The questions in the survey cover rich demographic information, a multitude of life perception issues, healthcare related information, as well as the intensity of social security ser-

[10]Granted, unraveling could well occur within the SHI system, since we must not forget that SHI is actually comprised of more than 100 companies that do have some minimal leeway in the manipulation of their risk pool. In fact, Bauhoff (2012) finds that SHI firms try to select customers on the basis of their geographic location. However, the argument here relates to the risk pool in the SHI system as a whole.

[11]It is important to note that this logic assumes no substantial changes in the pricing policies or contract space. A full welfare analysis would require predictions about the changes in the behavior of firms in response to any regulatory access changes.

[12]For for information on the SOEP panel please see http://panel.gsoep.de/

vices utilization. Among other things, the survey offers information on the type of health insurance coverage. As expected, most individuals report being insured in the SHI system. The number of individuals that were eligible for the PHI (excluding dependents) is 2,393 (13%) in the 2008 cross-section. 1,352 (56%) of these reported SHI enrollment and 1,041 (44%) were enrolled in the PHI. The survey further offers information on the self-reported insurance premiums and deductibles for the PHI insured, some information on other types of insurance policies in the household (e.g. life insurance), indicators of existing chronic diagnosis, BMI, self-reported information on the number of outpatient visits and hospital stays, as well as self-reported level of riskiness in different settings.

Table 3.9 summarizes the means of the key observed variables for all full-time employees and for all PHI-eligible full-time employees for years 2005-2009. The latter group is then split into those that reported SHI and those that reported PHI enrollment. The covariate means are compared for these two groups with a two-tailed t-test. We see that among all full-time employees, whose average monthly pre-tax income amounts to about 3000 EUR, only 8% of individuals report PHI enrollment. This share jumps dramatically to 35% among PHI-eligible employees. Given that PHI eligibility for employees is determined through income, it is not surprising that the PHI-eligible group reports much higher average income of about 5,800 EUR. This average is even higher for those employees that actually choose to enroll into the PHI - 6,150 EUR. Besides income, the main differences between the PHI-eligible and the all-employees groups are in covariates that we would expect to be highly correlated with income. Specifically, the PHI-eligible pool is older and more male. Perhaps unexpectedly, the PHI enrollees do not seem to be on average substantially different from the PHI-eligible SHI-enrollees, neither in demographic, nor in healthcare-related outcomes. In the next section I use econometric analysis to provide more detailed comparisons of these two groups.

3.3.2 Cream-skimming from the SHI system

The goal of this Section is to identify the patterns of any distinct sorting of individuals across the two insurance systems. As discussed in Section 3, ex ante the nature of net selection into the PHI system appears to be ambiguous. The ambiguity is especially stark if we believe that individuals choosing among the two horizontally differentiated insurance systems may have preferences for either one of the systems that are orthogonal to their expected healthcare expenditures. The key challenge for the empirical identification of selection between the systems is the need to disentangle the ex ante selection into the PHI system, from ex post causal effects of PHI enrollment. To address the identification challenge, the main analysis of this Section relies on the instrumental variables approach within a fuzzy regression discontinuity design. Specifically, I employ the income-based eligibility threshold below which individuals have to enroll into SHI, as an instrument for the individual's enrollment into the PHI. This allows me to identify the causal effect of the PHI enrollment and separate it from selection. In addition to the regression discontinuity analysis, I also discuss non-parametric evidence on the probability of chronic diagnosis by different types of insurance, as well as the development of covariates of the SHI-insured above and below the

PHI-eligibility threshold.

For the analysis of selection in this Section, I construct a data extract from the 2005-2009 SOEP waves. This extract consists of the individuals reporting to be working full-time, either as employees or self-employed. I thus exclude students, retirees and workers that are entitled to the civil servant medical benefits. I further trim the sample to include individuals whose monthly income is reported to be between 400 EUR and 50,589 EUR and exclude individuals that are younger than 25 or older than 65. The latter again allows me to ensure that there are no students and retirees in the sample, since students are often insured with their parents or student insurance, while retirees are locked into their respective insurance choices and no sorting can be taking place among them. The final sample includes 31,112 individual-year observations on about 6,000 individuals. The median monthly pre-tax income in the remaining sample is about 2,600 EUR/month. 90% of the sample lies in the monthly pre-tax income interval between 1,200 EUR and 7,000 EUR. About 14% of the individuals in each year report having a full private insurance coverage. Half of these are employees and the other half are self-employed. The income of the PHI-insured is about one standard deviation higher than the average income in the sample.

The final sample provides me with the information on the individual's basic demographic characteristics, such as age and gender, and the level of monthly income that the individual reports. I then observe whether the individual reported being in the PHI or the SHI system. System choices are very stable across years, so it is suitable to think about these choices using the static utility framework, which is employed in the next Section. At the same time this stability implies that any analysis of selection between the systems refers to the steady-state composition of the risks in the systems only, since the observations on switchers in their first year of switching are very scarce. I further observe several risk and taste-related characteristics of the individuals, such as the self-reported risk aversion, health status and satisfaction, sports affinity, healthy eating habits, smoking, body-mass index, disability and chronic diseases. Finally, I observe indicators of self-reported individual healthcare utilization, such as the number of physician office visits in the past three months and the number of hospital stays in the past year. Table 3.9 provides detailed summary statistics of the sample along different data cuts. Table 3.10 compares the covariate means within a 500 EUR bandwidth around the PHI eligibility cutoff. The samples on different sides of the cutoff appear to be similar on the covariates.

Correlation between healthcare utilization and PHI enrollment

To summarize the *prima facie* evidence on the relationship between the insurance system and the healthcare utilization, I use the following linear specification for the expected healthcare utilization outcomes as a function of the type of insurance.

$$E[Y^{outcome}|X,\, PHI] = \alpha PHI + \beta X$$

The outcome variables include the number of inpatient and outpatient visits, both unconditional

and conditional on having at least one visit, as well as the probability of having at least one visit. The parsimonious set of control covariates includes age, gender and income. I allow for a different intercept and a different slope for the self-employed individuals. Table 3.1 reports the OLS results for the regressions on the PHI-eligible sample. We observe that older individuals utilize more healthcare services, while women are more likely to visit a physician. Having PHI is correlated with a lower likelihood of a physician visit for the self-employed, while the effect on the outpatient behavior of the employees is negative, but not statistically significant different from zero. The connection between insurance status and inpatient stays of the employees is more precise than for outpatient visits. Specifically, it appears that employees are likely to have 0.27 fewer stays in the hospital per year, conditional on having at least one stay, if they have PHI insurance.

The association between the PHI enrollment and the utilization of healthcare observed in the linear specification warrants further analysis, as the choice of the PHI by an individual may be correlated with the unobserved characteristics that also determine the level of healthcare utilization. In particular, the OLS coefficients contain a mixed effect of selection into the PHI and the causal effect of the PHI on the individual's healthcare utilization. Even though the equation conditions on the key determinants of healthcare consumption such as age and gender, there may be taste characteristics of individuals that both induce the choice of the PHI and lesser or more significant healthcare utilization. For instance, if an individual likes to go to a physician a lot (for instance, for preventive care), then this individual may choose to buy PHI that provides better experience in the healthcare system and go to a physician more because of this idiosyncratic taste, even if conditional on age, gender, and income one would not predict higher utilization. This would be an example of selection. At the same time, the PHI provides different cost-sharing mechanisms than the SHI and thus we would expect less moral hazard in the PHI system.

Therefore, in order to more accurately characterize the nature of selection between the SHI and the PHI systems, I need to disentangle the causal effects of having the PHI on the level of healthcare utilization from the ex ante selection. Before proceeding with the causality analysis, it is important to emphasize that the causal or "moral hazard" and selection effects in this setting themselves include a multitude of potentially countervailing forces. Specifically, the selection effects here are a combination of any strategic customer screening by the PHI firms and individual preferences that lead different individuals to apply for and accept PHI contracts. The causal effect of the PHI may include the classical moral hazard argument, according to which the higher degree of cost-sharing should decrease the demand for healthcare. At the same time, PHI causality could also include the physician-induced demand argument, whereby the physicians, whose remuneration is substantially higher under the PHI, induce more demand from patients. A countervailing causal force would exist if PHI-insured are treated better and thus need less healthcare service. Lastly, if PHI patients face shorter waiting times and more convenient service, they could be inclined to more utilization of healthcare. The available data does not allow me to disentangle any of these forces separately; therefore, it is useful to keep in mind that my empirical findings of selection and causality will necessarily reflect the net of all these different influence channels.

Regression discontinuity analysis of the PHI's causal effect on healthcare utilization

To identify the causal effects of the PHI and subsequently separate them from selection, I need an instrumental variable that would be highly predictive of whether an individual has a PHI or not, but at the same time not related to the unobserved characteristics that may influence both the individual's utilization of healthcare and the choice of PHI conditional on the observed covariates. I exploit the regulatory break in PHI eligibility as an instrument for PHI enrollment. In the German system, the access to the choice between the SHI and the PHI for employees is restricted by the government. Only employees whose income crosses an annually set eligibility threshold may choose to opt out of the SHI system into the PHI. If the income eligibility boundary is binding, we would expect that there is a discontinuity in the probability of enrolling into the PHI at the income eligibility cutoff. In other words,

$$Pr(PHI_i = 1) = \begin{cases} g_1(income_i) & \text{if } income_i \geq cutoff \\ g_2(income_i) & \text{if } income_i < cutoff \end{cases}, \text{ where } g_1(\cdot) \neq g_2(\cdot)$$

This setting corresponds to a fuzzy regression discontinuity design, where I use the discontinuity in the probability of treatment as the instrument for the treatment status.[13] The discontinuity design is fuzzy, since the crossing of the eligibility threshold only gives the individual a choice to take up the PHI treatment, rather than imposing a switch to the PHI. The key identifying assumption in this setting is that individuals cannot precisely manipulate on which side of the cutoff they are, to gain the treatment. Given that income is likely reported with measurement errors (especially with rounding) and that employers tend to set rounded wages or use the insurance income cutoff as a wage benchmark, any statistical evidence of no manipulation should be interpreted with care. Nevertheless, histogram analysis of the density of observations at different levels of income suggests no evidence of heaping at the cutoff. Naturally, this still does not preclude the possibility of manipulation cases at the threshold. Given the employer's awareness of the threshold, it seems plausible, however, that employers would offer salaries that are right at the eligibility threshold or sizably lower. In this case, individuals are unlikely to have sufficient bargaining power to negotiate their wages up in the latter case, but at the same time it is also unlikely that individuals would choose their employment based on the PHI eligibility in the former case.

I start with the estimation of the first stage regression that should confirm that there is a strong relationship between the instrument and the endogenous decision to enroll in the PHI. I use different levels of polynomial controls that are allowed to differ above and below the cutoff and I center the income running variable at the cutoff. The latter makes the interpretation of results easier and allows me to combine observations from different years that had different cutoff levels.

[13]The fuzzy RD discussion here follows Angrist and Pischke (2009)

$$E[PHI|income] = \gamma_1 + \gamma_2 Above\, cutoff + \beta f(income - cutoff) + \delta f(income - cutoff) \times Above\, cutoff$$

Both the graphical representation of the probability of the PHI-enrollment as a function of income in Figure 3-4 and the regression coefficients in Table 3.2 confirm that PHI-eligibility induces individuals to take up the treatment after they cross the eligibility cutoff. As Figure 3-4 shows, there appears to be no relationship between the income cutoff and the PHI enrollment probability for the self-employed, for whom the income mandate doesn't apply. The simpler versions of the first stage regression without cutoff-centering yield similar results. The results remain quite consistent across different polynomial specifications. [14]

Having established the presence of the first-stage relationship, I proceed with the analysis of the reduced form specifications. As Angrist and Pischke (2009) point out, the analysis in this fuzzy regression discontinuity setting depends on whether we can document a discontinuity in the outcome variable at the cutoff conditional on the non-linearities of the forcing variable. Graphical representation of the reduced form in Figure 3-5 shows no evidence of a discontinuity between the health utilization outcomes at the income PHI-eligibility cutoff. The econometric reduced form specification centers the income variable at the cutoff and allows for different polynomials before and after the cutoff.

$$E[Y^{outcome}|X] = \alpha_1 + \alpha_2 Above\, cutoff + \beta f(income - cutoff) + \delta f(income - cutoff) \times Above\, cutoff$$

Table 3.3 summarizes the reduced form coefficients. The estimates are imprecise, but broadly confirm the intuition from the graphical evidence that there is no stark jump in the average utilization of the healthcare services below and above the PHI-eligibility cutoff. This suggests that there is no strong evidence of net causal effect of the PHI. While this may appear surprising in light of the cost-sharing arrangements in the PHI, one should keep in mind that this result does not imply that cost-sharing arrangements in the PHI do not have any effect on the consumption of care. The concept of moral hazard differences between the two insurance systems again consists of several potentially countervailing forces. First, we would expect cost-sharing arrangements to decrease the demand for care. At the same time, however, better and faster service towards the PHI-insured may have the opposite effect. Lastly, the high remuneration of physicians by the PHI providers may create incentives for more physician-induced demand for care that would again countervail the

[14] Considering the potential measurement error in income, local results around the cutoff may be misleading, since the observations around the cutoff may have been misclassified. Note that the graphical evidence suggests that there are a number of observations very close to the cutoff that have a fairly high probability of PHI enrollment, even if their income is reported to be below the eligibility level. The first reason for such observations may be a measurement error in income that leads me to misclassify the individual's eligibility. Secondly, the German health insurance regulation allows individuals that opted out to the PHI at some point and then their income dropped below the current eligibility threshold, to sign a waiver for the re-entry of the SHI.

decrease due to cost-sharing.

Table 3.4 summarizes the results of the first stage and the reduced form specifications in the 2SLS regression. Given the imprecision of the reduced form results, the 2SLS coefficients are also imprecise. In the OLS regressions at the beginning of this section, we saw that the only significant difference in the utilization of healthcare by the PHI-insured appears to occur in the frequency of stays at an inpatient facility conditional on having any stay in a given year. Since the 2SLS coefficients are not precise estimates, I cannot subtract them from the (also imprecise) OLS results to get the magnitude of the net selection between the PHI and the SHI systems. However, I can use the information from the confidence intervals of these two specifications to calculate the range of possible selection levels that I cannot reject at 5% confidence levels. Consider, for example, the total number of outpatient visits in three months as the outcome variable. The OLS results suggest that the effect of the PHI on this number lies between [-0.33, 0.07]. The 2SLS coefficient for the same outcome variable has a confidence interval of [-1.36, 0.0009]. Then, the data would suggest that we cannot reject that the causal effect of the PIII is to induce 1.36 fewer outpatient visits. If this causal effect were true, then to get the combined effect of causality and selection to be at least -0.33, we would need that individuals with on average one *more* outpatient visit get selected into the PHI. This would imply advantageous selection. If we consider the other side of the confidence interval that suggests zero causal effect of the PHI on physician visits, we would conclude that individuals with on average 0.33 fewer physician visits switch to the PHI. That is the highest level of adverse selection would be 0.33 outpatient visits per quarter. The econometric analysis also cannot reject that both the causal and the selection effects of the PHI on healthcare utilization are zero.

Additional evidence on selection between the PHI and the SHI

The "zero" effects appear to be supported by two additional pieces of evidence. First, Figure 3-1 shows no evidence of any difference in the probability of having a chronic diagnosis at different levels of income for individuals in different insurance systems. If one believes that these diagnosis cannot be strongly related to which kind of insurance the individual has (and this is plausible, since the level of SHI coverage is very high), then these graphs would suggest that PHI companies do not manage to differentially select individuals that will not have chronic/expensive illnesses. The only evidence of selection appears to happen on the diabetes-diagnosis. We observe that the probability of having a diabetes diagnosis is consistently lower for the PHI-insured with income above the eligibility threshold. These differences could be related to the nature of PHI pricing, as discussed in Section 3, whereby PHI companies observe the individuals at a fairly early age and thus may be unable to predict and screen on most of the diagnosis, except for diabetes.

Figures 3-2 and 3-3 provide further evidence for the no-selection hypothesis. These figures plot the development of a number of covariates, which may plausibly not contain a causal effect of the PHI, for the SHI-insured across the income eligibility threshold. The covariates include the average age, the fractions of older and younger individuals, the average BMI, health-related risk

aversion, the affinity to smoking, to healthy eating, sports and the fraction of disabled. If PHI disproportionately selects individuals with certain levels of these covariates, we would expect to see a jump in the level of the covariates for those that stay in the SHI insurance. To detrend the development of the variables from the relationship with income, I plot the residuals from the regression of the covariates on income using the sample to the left of the cutoff and calculating out-of-sample residuals to the right of the cutoff. Again, these graphs suggest no evidence of a discontinuous net selection at the income eligibility threshold. As discussed in Section 3, finding no evidence of selection would imply that at least within some income bandwidth, the PHI eligibility threshold does not serve the purpose of maintaining the SHI risk portfolio. This conclusion renders the welfare analysis of the next section relevant for the policy discussion about the necessity or the level of such market access regulation.

3.4 Preferences for private insurance and the welfare effects of the SHI mandate policy

3.4.1 Premium payments for SHI and PHI

The key ingredient in the individual's choice between the statutory and the private systems is the relative price that the individual faces on both markets. Since I can only observe the prices for the PHI system if the individual chose to enroll into the PHI and reported the monthly premiums, I have to simulate the PHI premiums for all the SHI-enrollees. Together with the calculation of the SHI premiums, I can then derive hypothetical incremental difference between the premiums. Since the calculation of the counterfactual SHI price is straightforward, as it is set by law, I start with the discussion of the PHI prices.

To simulate the private insurance pricing for the individuals that chose to stay with the SHI, I use the prices reported by the individuals that did choose PHI and project these prices on the observable characteristics of these individuals. I consider five years of the data, years 2005-2009, where PHI pricing feasibly followed the same broad principles.[15] The sample for five years of the PHI prices has a total of 8,429 individual-year observations on about 1,500 individuals. The data includes observations on the monthly premium paid into the PHI, the number of people covered under this policy, whether the individual was self-employed or a civil servant, age, gender, BMI, number of outpatient visits in the past three months of the survey year, number of inpatient visits, whether the individual is a smoker, whether the individual experienced prolonged work disability, and indicators for diabetes, asthma, cardiac conditions, cancer, stroke, migraine, high blood pressure, depression, dementia, other illnesses, whether no serious diagnosis has been made.

[15] Although using several years of the panel gives me the advantage of having a larger sample size, there is also a disadvantage in using these years. In the period between 2007 and 2009, there was a short regulatory change that required the individuals to cross the income eligibility threshold three years in a row before they were eligible to switch to the PHI. However, if anything, this feature would lead me to underestimate the demand for the PHI and thus provides a more conservative specification.

Private health insurance companies use annuity pricing formulas that include information on the individual's age, gender, and health to arrive at the monthly premium payments for each individual. In theory these premiums correspond to the average monthly payment required to cover expected life-time medical expenses, given the levels of morbidity and mortality and should remain stable over time. In practice, however, the payments have to be continuously re-adjusted to rising medical costs and increasing life expectancy. Given the details of the pricing system, individuals that have been in the PHI for longer and are older usually get higher annual premium increases than younger individuals. Since I do not observe the individual PHI entry age, I use an unrestricted linear specification to approximate the expected conditional price that individuals would face for a given age, gender and health status.

The general econometric specification is as follows:

$$E[P_i^{PHI}|\theta_i] = \alpha_0 + \alpha_1 \cdot female_i + \alpha_2 \cdot age_i + \beta \cdot healthcare\,indicators_i + \gamma \cdot employment\,type\,fixed\,effects_i$$

where the P^{PHI} refers to the reported prices in EUR for the years 2005-2009, $female$ equals one if the observation is for a woman, age has age in years, $healthcare\,indicators$ refer to different controls for chronic conditions or healthcare utilization. Employment fixed effects capture the differences in the premiums that self-employed, civil servants and full-time employees pay. These effects are included in up to third-level interactions to account for potentially differential prices for, for instance, full and part-time employees. Table 3.5 summarizes the key coefficients for different specifications.

The regression implies that a 40-year old male that is a full-time employee would pay on about 400 EUR a month for his private health insurance. Reformulating the regression in percentage terms by doing a log-transformation of the price variable, we get that each additional year increases the premium by about 2%. It is clear that the pricing is also sensitive to the individual's gender. It appears that women have both a different level and a different age slope in the PHI pricing. This is not surprising, since reported costs for women in the PHI system indeed have different levels in different age groups. Single individuals pay less for their insurance, which is natural, since they on average pay for fewer policies. The positive coefficient on income most probably captures the fact that individuals with higher income choose more comprehensive, and thus more expensive, coverage. The linear approximation of the PHI pricing rules accounts for a substantial amount of variation in the prices, with R^2 equal to 0.42 in the log-specification with diagnosis-specific controls.

Using this approximation of the average PHI prices, I calculate the counterfactual PHI premiums for the PHI-eligible individuals that had chosen SHI in the years 2005-2009. This counterfactual calculation is based on non-trivial assumptions. The first assumption is that conditional on the same demographics, individuals that didn't change to the PHI would have changed at the same age and would have chosen a similar PHI plan as the individuals that did change. Second, conditional on the same demographics, individuals that didn't switch would have been in a similar health

condition at the time of the switch. In short, I assume that individuals that didn't switch to the PHI would have faced the same pricing rules as the individuals that did switch. Given the discussion in Sections 3 and 4 that points out to the fact that even though PHI firms are allowed to price-discriminate on the observables at the time of enrollment, their ability to actually distinguish healthcare expenditure risks of individuals of the same age and gender (where these individuals are mostly middle-aged high earners) may be limited. In that sense, the assumption that individuals would have faced the same prices at the time of the application is less stringent.

The left panel of Figure 3-6 compares the densities of predicted and observed prices for individuals that reported prices. The right panel compares the reported PHI-premiums with out-of-sample predictions for the currently SHI-insured. The lower part of the panel provides the same comparison for the full-time employees only. The densities in Figure 3-7 show the difference between the observed prices for PHI-insured and the predicted prices for the PHI-eligible SHI-insured. These plots indicate that the pricing regression performs well in predicting the total in-sample premium levels, while it performs substantially worse in the predictions for the employees only. Further, densities indicate that PHI-eligible employees chose policies with higher monthly premiums than the self-employed. This is not surprising, since unlike the self-employed that are PHI-eligible at any income level, the PHI-eligible employees are high earners that could be expected to choose more expensive plans.

In the next step, I apply the rules for SHI-premium calculations from respective years to get imputed SHI-premiums for both the SHI and the PHI-insured. These premiums then allow me to calculate the hypothetical price difference that the PHI-eligible individuals face. The predicted SHI premiums and the distribution of premium differences is illustrated in Figure 3-8. Note that the SHI premiums have two different levels. The lower level corresponds to the premiums paid by the employees, for whom the employer pays about half of the premium, while the upper level are the monthly premiums paid by the self-employed. The right panel of the figure plots imputed premium differences between the PHI and the SHI for employees.[16] For the majority of individuals, PHI appears to be more expensive than the SHI. At the same time, it appears that there is no much difference in the predicted premiums for those that actually chose PHI and those that decided not to opt out of the SHI, which would suggest non-pricing preferences to be important for the demand decisions.

[16]Note that this calculation faces a certain data-interpretation problem. In the German system, the employees can apply their employer's health insurance subsidy to their private insurance coverage. However, in the survey, when the employees are asked about their PHI premiums, the question does not specify whether the respondents should specify the monthly premium after or before the employer's subsidy. To remain conservative, I assume that the individuals reported the post-subsidy premium that they actually pay. Therefore, when calculating the predicted differences in the SHI and PHI premiums, I also apply the post-subsidy levels of the SHI premiums. If the assumption about the data is wrong and individuals reported the full premium rather than the post-subsidy premium, then I would be dramatically overestimating the PHI prices and thus underestimating the probability that the individuals switch to the PHI and the implied welfare effects.

3.4.2 Preferences for PHI

To formulate a simple model of demand for the private insurance, I go back to the consumer utility characterization introduced in Section 3. This utility characterization assumes that individual's utility from insurance is the difference between the individual's willingness to pay and the premium that the individual faces. That is, the utility can be characterized by the following:

$$U_{ij} = u_j(\theta_i) - p_j(\theta_i) \tag{3.1}$$

This representation of the individual's utility is readily mapped into a standard discrete choice framework.[17] In this framework, we would like to specify the observed and the unobserved (to the econometrician) portions of utility. That is, we want to decompose utility into the representative utility V_{ij} that is the observable part of utility and the unobservable part ϵ_{ij}. The unobservable part is defined as a residual in such a way that $U_{ij} = V_{ij} + \epsilon_{ij}$. Comparing this to Equation 3.1, we see that the unobserved part of utility will be the unobserved part of the willingness to pay $u_j(\theta_i)$ and the representative utility will include the observed part of the willingness to pay as well as prices.

As in a standard discrete choice framework, I assume that the observed part of utility is a function of the alternative's attributes Z_{ij} that may vary by the decision maker, but at least vary by the alternative, as well as some attributes of the decision maker X_i that are the same across the alternatives. Then, the observed part of utility can be approximated linearly as:

$$V_{ij} = \beta X_i + \gamma Z_{ij}$$

Specifically, suppose each individual can be characterized by a set of attributes X_i that doesn't change depending on the choice of insurance, but does affect individual's preference for one or the other choice. In out setting X_i consists primarily of demographic factors, such as age, gender, employee or self-employment status, income, sport affinity, smoking, BMI, and the degree of risk aversion.[18] The observed choice-varying attributes in our setting only include the different prices that individuals face for different insurance choices, i.e. $Z_{ij} \equiv p_{ij}$.

I explicitly include the alternative-specific fixed effect into the utility model. This constant captures the average choice-specific valuation of the unobserved portions of utility and thus automatically normalizes the mean of the unobserved utility ϵ_{ij} to zero. The final utility model then takes the following form:

$$U_{iPHI} = \beta_{PHI} X_i + \gamma p_{iPHI} + \xi_{PHI} + \epsilon_{iPHI}$$

[17] In the formulation of the discrete choice model I closely follow the treatment in Train (2003).

[18] Arguably factors such as sport affinity, smoking and BMI may be endogenous to the insurance choice. For example, private insurance may work on encouraging the insured to quit smoking or loose weight. However, I abstract from this possibility for now and assume that individuals don't explicitly and differentially change their traits because of different insurance choices.

$$U_{iSHI} = \beta_{SHI} X_i + \gamma p_{iSHI} + \xi_{SHI} + \epsilon_{iSHI}$$

The individual is assumed to choose the alternative that maximizes utility. The probability that an individual i chooses PHI over SHI is then given by

$$\begin{aligned} Pr(U_{iPHI} > U_{iSHI}) = \\ =Pr(\beta_{PHI} X_i + \gamma p_{iPHI} + \xi_{PHI} + \epsilon_{iPHI} > \beta_{SHI} X_i + \gamma p_{iSHI} + \xi_{SHI} + \epsilon_{iSHI}) = \\ =Pr(\epsilon_{iSHI} - \epsilon_{iPHI} < \gamma p_{iPHI} - \gamma p_{iSHI} + \xi_{PHI} - \xi_{SHI} + \beta_{PHI} X_i - \beta_{SHI} X_i) \end{aligned}$$

Since utility is ordinal, the coefficients in this model are identified only up to a constant, so I normalize ξ_{SHI} and β_{SHI} to be zero. Then we get:

$$\begin{aligned} Pr(i\, choose\, PHI) = \\ =Pr(\epsilon_{iSHI} - \epsilon_{iPHI} < \gamma p_{iPHI} - \gamma p_{iSHI} + \xi + \beta X_i) \end{aligned}$$

That is, under the assumption of separability of valuation and price, demand for PHI depends on the difference in the valuation of the products and the corresponding difference in price as well as the density function F of the difference in the error terms $F(\epsilon_{iSHI} - \epsilon_{iPHI}) \equiv F(\tilde{\epsilon_i})$. In other words:

$$\begin{aligned} Pr(i\, choose\, PHI) = \\ = \int I(\tilde{\epsilon_i} < \gamma p_{iPHI} - \gamma p_{iSHI} + \xi + \beta X_i) dF(\tilde{\epsilon_i}) \end{aligned}$$

Assuming extreme value Type 1 distribution on the error terms, which implies that $F(\tilde{\epsilon_i})$ is logistic, we get a standard logit[19] closed-form representation:

$$Pr(i\, choose\, PHI) = \frac{e^{\beta X_i + \gamma p_{iPHI} + \xi}}{e^{\beta X_i + \gamma p_{iPHI} + \xi} + e^{\gamma p_{iSHI}}}$$

which is equivalent to:

$$\frac{(e^{\beta X_i + \gamma p_{iPHI} + \xi}) e^{-\gamma p_{iSHI}}}{(e^{\beta X_i + \gamma p_{iPHI} + \xi} + e^{\gamma p_{iSHI}}) e^{-\gamma p_{iSHI}}} = \frac{e^{\beta X_i + \gamma (p_{iPHI} - p_{iSHI}) + \xi}}{1 + e^{\beta X_i + \gamma (p_{iPHI} - p_{iSHI}) + \xi}}$$

[19]Note that in the binary choice situation that we are analyzing, the Independence of Irrelevant Alternative property that makes logit model unattractive for differentiated goods analysis is not applicable. However, the inability of logit to capture the correlation in unobserved parts of utility over time is a fairly significant drawback for the choice of insurance and thus using a less restrictive model would be a useful extension of the paper. Furthermore, an extension to a more flexible mixed logit specifications would also help to capture preference heterogeneity more accurately.

The latter can be readily used to formulate and estimate the standard logit maximum likelihood. Table 3.6 reports the marginal effects from the logit estimates as well as linear-probability model comparison for different specifications of the representative utility. The sample in this estimation is restricted to non-civil servant, full-time employed individuals that report having voluntary SHI insurance or PHI insurance. In other words, to those individuals, who appear to have a clear choice between the PHI and the SHI systems. To better understand the differences in demand across different quantiles of the income distribution, I also present the preferred demand specification (columns (3) and (4) of the full-sample specification) for three separate income brackets in Table 3.7.

All specifications suggest that demand for the PHI is downward-sloping in price. Conditional on prices, there is some taste-heterogeneity in the preferences for the insurance system. Thus, older, self-employed, higher-income, and healthy-eating affine individuals as well as smokers are more likely to get the private health insurance conditional on the prices. At the same time, there appears to be no heterogeneity of tastes based on gender, BMI and risk aversion. Individuals with disability are less likely to enroll into the PHI - this may be a reflection of both tastes (e.g. aversion towards deductibles) as well as rejections by the PHI companies. Interesting taste heterogeneity is captured by the indicators for whether individuals employ household help and whether they have strong center-left political views. Supposing that individuals with household help highly value convenience, the strong effect of this coefficient would confirm the presence of convenience preferences in the choice of the horizontally differentiated insurance plans. The elasticity of demand towards the price differential conditional on the demographics appears to increase with income levels. At the same time the residual impact of the "convenience" preferences, which are proxied by whether or not the individual employs household help, fades at higher income levels. In the next Section, I use demand estimates to predict counterfactual PHI demand for individuals with income below the income eligibility threshold.

3.4.3 Counterfactual welfare effects of SHI mandate

In order to quantify the welfare effects of the income eligibility threshold on the individual utility, I need to make predictions about the counterfactual utility levels. The demand estimation in Section 5.2 provides me with an approximation for the "observed" or the representative part of utility for the PHI relative to the SHI as a function of observable covariates. In the demand estimation, I can normalize the observable utility for the SHI to zero, using the relative price for the PHI rather than absolute. Consequently, the utility for the PHI just gives the additional representative utility relative to the SHI. Therefore, I can use the demand regression coefficients to construct empirical difference in the representative utility:

$$\hat{V}_{iPHI-SHI} = \hat{\beta}_{PHI-SHI} X_i + \hat{\gamma} p_{iPHI-SHI} + \hat{\xi}_{PHI-SHI}$$

The change in the expected consumer surplus from the scenario with the PHI choice available

versus without the choice is then found using the standard consumer surplus formulation in the logit framework (Train, 2003) :

$$\triangle E(CS_i) = \frac{1}{\hat{\gamma}}[ln(1 + e^{\hat{V}_{iPHI-SHI}})]$$

Since the coefficient on price gives the marginal value of money, dividing the consumer surplus equation by this coefficient allows me to interpret the utility difference in monetary terms. The histograms in Figure 3-9 show the spread and levels of the representative utility calculations for the PHI-eligible sample of employees on which the demand estimation was done. The "observable" portion of the willingness to pay for the PHI across almost all models lies between -400 EUR and +400 EUR per capita. Across all models less than 50% of the distribution lies in the positive region. That is, more than half of the individuals in the sample would have higher representative utility from the SHI coverage. This, however, does not take into account the expectation of the unobserved part of the random utility model.

To calculate the partial equilibrium welfare impact of the PHI eligibility restriction, I compare the expected consumer surplus with and without the PHI option for the currently ineligible population. The welfare calculations in Table 3.8 provide the implied surplus loss per capita as well as the implied total surplus loss for the population, where the latter is computed using sample weights for the survey. The calculations imply that removing the PHI access restriction completely would create monthly gains in surplus per capita of 4-39 EUR for currently ineligible individuals. Looking at the counterfactual policy of moving the threshold by 500 EUR rather than completely removing it, gives us a better idea of surplus gains for the group of individuals that are actually likely to choose PHI. Specifically, over this population the per capita monthly surplus gains are on the order of 15-60 EUR. This scenario also implies a gain of 75-200 mn EUR for the population, calculated using sample survey weights.

Figure 3-10 plots the differences in the expected consumer surplus gains that is predicted for different income levels. As expected, lower income is correlated with low to none surplus gains from the extra choice. Figure 3-11 shows how changes in the expected consumer surplus develop over different levels of policy changes. In accordance with the observation of decreasing gains in the previous Figure, we see that welfare gains flatten out the farther away we move the restrictive income threshold. At the same time the calculations far out of sample should be interpreted with caution. Individuals with substantially lower income levels may be systematically different in their preferences from the high earners, in which case the extrapolation of preferences out of sample cannot be too informative. To emphasize the focus on the observations with income levels closer to the current income threshold, the welfare function Figures provide a more detailed picture of the predictions within 1000 EUR bandwidth of the threshold, where the predictions of different demand models are also much more congruent.

It is important to note here that these welfare impact calculations are only partial equilibrium results, since I keep the pricing in both systems fixed and only introduce the counterfactual changes in the access restriction policy. Furthermore, these welfare calculations consider only the impact on

the individuals and do not take into account any effects on the insurance providers. Ideally, with a richer model and richer data, one could conduct a full welfare analysis that would take into account the changes in the composition of enrollees in both the PHI and the SHI with a policy change, how that would impact prices and how those would in turn impact choices. Note that prices in the SHI could change with a change in the composition of enrollees even if there is little adverse or advantageous selection into the PHI as I seem to observe in Section 4. This is due to the fact that SHI covers a lot of enrollees at the bottom of the income distribution and individuals without any income, so any movements outside of the SHI could result in price adjustments. To reiterate, for the welfare calculations in Table 3.8, I assume that the only piece of the system changing is the eligibility threshold. Under this assumption, restricting access to the PHI system appears to create a sizable welfare loss for individuals below the eligibility threshold.

3.5 Conclusion

In this paper I have analyzed the role of taste preferences, risk selection, and mandated enrollment regulation for the market of two competing systems within the institutional setting of two-tier German health insurance. I find clear evidence for the presence of taste preferences in the demand for private health insurance, which offers more convenience and higher service quality, but also higher levels of cost-sharing. Since it appears plausible that the identified taste heterogeneity is not directly related to expected healthcare expenditures, this finding is consistent with the idea that there is scope for horizontal differentiation of health insurance products that is valuable for consumers. Since the German government mandates enrollment into the social insurance system for employees with income below a certain threshold, I have also addressed the natural question about the magnitude of welfare implications of this policy. Indeed, I find that this regulation is binding and creates a substantial *ceteris paribus* consumer welfare loss. The relevance of these welfare calculations depends critically on the nature of selection between the two insurance systems. In my analysis of selection, I find no strong evidence of adverse risk sorting from the SHI to the PHI. While this finding may appear surprising, I argue that the annuity nature of the PHI premiums and the regulatory policy of life-time health insurance contracts could explain the limited degree of screening possibilities in the system. A more precise analysis of selection is critical for the policy implications, however, since if there were adverse selection into the PHI system, the regulatory access restriction would play an important role in the maintenance of a sustainable risk pool in the SHI system that would weigh against individual welfare losses from the regulation.

Table 3.1: Healthcare utilization and insurance type

	(1) OLS Physician visits	(2) OLS P(Phys visit>=1)	(3) OLS Physician visits	(4) OLS Hospital stays	(5) OLS P(Hosp stay>=1)	(6) OLS Hospital stays
PHI	-0.131	-0.0339	-0.0553	-0.0237*	-0.00395	-0.268***
	[-0.330,0.0692]	[-0.0693,0.00141]	[-0.324,0.214]	[-0.0470,-0.000403]	[-0.0202,0.0123]	[-0.427,-0.110]
Age	0.0239***	0.00421***	0.0215***	0.00260***	0.00185***	0.00367
	[0.0155,0.0323]	[0.00274,0.00568]	[0.00924,0.0339]	[0.00162,0.00358]	[0.00117,0.00254]	[-0.00181,0.00915]
Female	0.398***	0.114***	0.133	0.00237	0.00420	-0.0168
	[0.205,0.591]	[0.0818,0.147]	[-0.132,0.398]	[-0.0184,0.0231]	[-0.00996,0.0184]	[-0.189,0.155]
log Income	-0.0474	-0.00636	-0.0528	0.00400	0.00460	-0.00696
	[-0.172,0.0770]	[-0.0303,0.0176]	[-0.246,0.140]	[-0.00972,0.0177]	[-0.00547,0.0147]	[-0.106,0.0923]
Self-employment	-0.176	-0.0848***	0.108	-0.0138	0.000290	-0.183*
	[-0.423,0.0711]	[-0.123,-0.0463]	[-0.255,0.472]	[-0.0403,0.0128]	[-0.0174,0.0180]	[-0.341,-0.0237]
Self-employed x PHI	-0.227	-0.0564*	-0.131	0.00783	-0.0106	0.293**
	[-0.533,0.0792]	[-0.109,-0.00362]	[-0.590,0.328]	[-0.0241,0.0398]	[-0.0334,0.0122]	[0.0846,0.501]
Constant	1.003	0.465***	2.252**	-0.0552	-0.0538	1.264**
	[-0.0209,2.027]	[0.260,0.669]	[0.666,3.839]	[-0.170,0.0598]	[-0.137,0.0298]	[0.336,2.192]
Observations	9485	9485	5353	9472	9472	641

95% confidence intervals in brackets
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table provides descriptive statistics for the relationship between the utilization of healthcare and demographic factors as well as the insurance policy. The sample includes all working non-civil servant individuals in 2005-2009 that would have been eligible to buy the private insurance. That is, the sample includes the self-employed and the employees with income above the regulatory threshold. Six outcome variables are: 1) the total reported number of outpatient visits in the three months before the survey date; 2) the probability of reporting any outpatient visit in this time period; 3) the number of outpatient visits conditional on having had at least one; 4-6) the same for inpatient visits within a year from the survey date. Note that the utilization information only includes the number of visits (so no cost information is available) and it is self-reported by the survey respondents. Since the coefficients on the insurance type are not precise, 95% confidence interval is reported for an easier interpretation of the possible effects. The regressions allow for a different slope and a different intercept for the self-employed individuals. Standard errors are clustered at the individual level.

Table 3.2: First stage regressions: relationship between PHI eligibility and PHI enrollment

	(1) OLS PHI	(2) OLS PHI	(3) OLS PHI	(4) OLS PHI	(5) OLS PHI
Above cutoff	0.246***	0.175***	0.121***	0.0825***	0.0821***
	(0.0152)	(0.0176)	(0.0203)	(0.0227)	(0.0227)
Cutoff deviation	0.0000195***	0.0000839***	0.000147***	0.000256***	0.000256***
	(0.00000221)	(0.00000972)	(0.0000246)	(0.0000476)	(0.0000476)
Deviation x above	0.0000129*	-0.0000264	-0.0000502	-0.000130*	-0.000128*
	(0.00000579)	(0.0000139)	(0.0000296)	(0.0000534)	(0.0000533)
Deviation squared		2.03e-08***	6.74e-08***	0.000000206***	0.000000205***
		(2.60e-09)	(1.47e-08)	(4.72e-08)	(4.72e-08)
Deviation squared x above		-2.20e-08***	-7.59e-08***	-0.000000222***	-0.000000221***
		(2.63e-09)	(1.47e-08)	(4.71e-08)	(4.71e-08)
Deviation cubed			9.68e-12***	7.28e-11***	7.27e-11***
			(2.72e-12)	(1.88e-11)	(1.88e-11)
Deviation cubed x above			-9.47e-12***	-7.19e-11***	-7.18e-11***
			(2.72e-12)	(1.88e-11)	(1.88e-11)
Deviation fourth				9.41e-15***	9.44e-15***
				(2.60e-15)	(2.60e-15)
Deviation fourth x above				-9.42e-15***	-9.45e-15***
				(2.60e-15)	(2.60e-15)
Age					-0.000566*
					(0.000246)
Female					0.0119*
					(0.00501)
Health satisfaction					0.00349**
					(0.00109)
Constant	0.0472***	0.0865***	0.107***	0.127***	0.125***
	(0.00448)	(0.00858)	(0.0126)	(0.0162)	(0.0215)
Observations	26828	26828	26828	26828	26788

Individual-clustered standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table provides several specifications of the first stage regression. The sample includes all individuals in 2005-2009 that are full-time employees, except for civil servants. The goal of the specifications is to determine whether crossing the PHI eligibility threshold increases the probability of taking-up a PHI policy. The specifications control for several polynomial trends of the explanatory variable - income expressed as deviations from the cutoff value - since the effect of crossing the threshold has to be observed beyond any non-linear relationship between the PHI choice and income. Adding higher-order polynomials dampens the effect of the threshold. The last specification also controls for the key demographics which, however, do not affect the level of change in the treatment probability at the cutoff. Crossing the threshold clearly increases the probability of enrollment into the PHI, but the treatment is not deterministic, which suggests a fuzzy regression discontinuity design. Standard errors are clustered at the individual level.

Table 3.3: Reduced form regressions: relationship between healthcare utilization and PHI eligibility

	(1) OLS Physician visits	(2) OLS P(Physician visit>=1)	(3) OLS Physician visits	(4) OLS Hospital stays	(5) OLS P(Hospital stay>=1)	(6) OLS Hospital stays
Above cutoff	-0.120	-0.0104	-0.132	0.00213	0.00178	-0.00556
	[-0.276,0.0352]	[-0.0371,0.0164]	[-0.339,0.0747]	[-0.0176,0.0219]	[-0.0110,0.0146]	[-0.177,0.166]
Cutoff deviation	0.0000649*	0.0000246***	-0.0000101	-0.00000586	-0.00000400	-0.00000830
	[0.00000343,0.000126]	[0.0000139,0.0000353]	[-0.0000920,0.0000718]	[-0.0000135,0.00000177]	[-0.00000907,0.00000108]	[-0.0000636,0.0000470]
Cutoff deviation x cutoff	-0.000104**	-0.0000327***	-0.0000202	0.00000284	0.00000178	0.00000688
	[-0.000173,-0.0000355]	[-0.0000454,-0.0000200]	[-0.000112,0.0000718]	[-0.00000566,0.0000113]	[-0.00000384,0.00000740]	[-0.0000768,0.0000906]
Age	0.0256***	0.00381***	0.0233***	0.00285***	0.00207***	0.00222
	[0.0208,0.0303]	[0.00306,0.00457]	[0.0172,0.0295]	[0.00226,0.00345]	[0.00169,0.00245]	[-0.00171,0.00614]
Female	0.676***	0.161***	0.320***	0.00945	0.00938*	-0.0335
	[0.573,0.779]	[0.145,0.177]	[0.192,0.448]	[-0.00311,0.0220]	[0.00154,0.0172]	[-0.129,0.0624]
Constant	0.623***	0.435***	1.849***	-0.0384**	-0.0239*	1.195***
	[0.398,0.848]	[0.396,0.473]	[1.553,2.144]	[-0.0666,-0.0101]	[-0.0424,-0.00537]	[0.983,1.408]
Observations	26828	26828	16504	26780	26780	1955
R^2	0.019	0.029	0.007	0.004	0.006	0.001

95% confidence intervals in brackets
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table provides linear specifications for the reduced form relationship between the utilization of healthcare and the eligibility for the PHI. The sample includes all individuals in 2005-2009 that are full-time employees, except for civil servants. The goal of the specification is to quantify any discontinuity in the outcome variable at the eligibility threshold beyond possible non-linearities in the model. The linear specification is reported here; however, higher order polynomials produce a similar results. In no specification is there statistically significant evidence of a discontinuity at the threshold. The six outcome variables are: 1) the total reported number of outpatient visits in the three months before the survey date; 2) the probability of reporting any outpatient visit in this time period; 3) the number of outpatient visits conditional on having had at least one; 4-6) the same for inpatient visits within a year from the survey date. Note that the utilization information only includes the number of visits (so no cost information is available) and it is self-reported by the survey respondents. Since the coefficients on the PHI eligibility indicator are not precise, 95% confidence interval is reported for an easier interpretation of the possible effects. Standard errors are clustered at the individual level.

Table 3.4: Instrumental variable specification for the fuzzy Regression Discontinuity design

	(1) IV Physician visits	(2) IV P(Physician visit>=1)	(3) IV Physician visits	(4) IV Hospital stays	(5) IV P(Hospital stay>=1)	(6) IV Hospital stays
PHI	-0.679	-0.104	-0.600	0.0139	0.0105	-0.00738
	[-1.359,0.000939]	[-0.225,0.0165]	[-1.551,0.351]	[-0.0713,0.0992]	[-0.0452,0.0661]	[-0.823,0.808]
Deviation x PHI	-0.000158**	-0.0000514***	-0.0000223	0.00000438	0.00000271	0.0000105
	[-0.000271,-0.0000446]	[-0.0000759,-0.0000270]	[-0.000165,0.000120]	[-0.00000882,0.0000176]	[-0.00000595,0.0000114]	[-0.000119,0.000140]
Female	0.688***	0.164***	0.327***	0.00916	0.00918*	-0.0335
	[0.584,0.792]	[0.147,0.180]	[0.197,0.456]	[-0.00359,0.0219]	[0.00127,0.0171]	[-0.135,0. 0681]
Age	0.0248***	0.00363***	0.0230***	0.00287***	0.00208***	0.00223
	[0.0200,0.0296]	[0.00286,0.00440]	[0.0169,0.0292]	[0.00227,0.00347]	[0.00170,0.00247]	[-0.00173,0.00619]
Cutoff deviation	0.0000807*	0.0000272***	0.00000365	-0.00000620	-0.00000424	-0.00000815
	[0.00000756,0.000154]	[0.0000145,0.0000400]	[-0.0000954,0.000103]	[-0.0000153,0.00000287]	[-0.0000103,0.00000180]	[-0.0000744,0.0000581]
Constant	0.685***	0.447***	1.890***	-0.0397*	-0.0249*	1.195***
	[0.433,0.937]	[0.404,0.490]	[1.561,2.220]	[-0.0714,-0.00804]	[-0.0457,-0.00408]	[0.965,1.425]
Observations	26828	26828	16504	26780	26780	1955

95% confidence intervals in brackets
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The table provides instrumental variables specifications for the relationship between the utilization of healthcare and the PHI. The PHI enrollment and the enrollment interacted with the running variable are instrumented by PHI eligibility and PHI eligibility interacted with the running variable. The sample includes all individuals in 2005-2009 that are full-time employees, except for civil servants. The goal of the specification is to quantify the causal effect of the PHI on healthcare utilization (i.e. a compounded moral hazard effect). A linear specification is reported here; however, higher order polynomials produce similar results or do not properly invert. In no specification is there a statistically significant strong causal effect of the PHI on utilization. The six outcome variables are: 1) the total reported number of outpatient visits in the three months before the survey date; 2) the probability of reporting any outpatient visit in this time period; 3) the number of outpatient visits conditional on having had at least one; 4-6) the same for inpatient visits within a year from the survey date. Note that the utilization information only includes the number of visits (so no cost information is available) and it is self-reported by the survey respondents. Since the coefficients on the PHI enrollment indicator are not precise, 95% confidence interval is reported for an easier interpretation of the possible effects. Standard errors are clustered at the individual level.

Table 3.5: Linear approximation to PHI pricing

	(1) OLS PHI premium	(2) OLS PHI premium	(3) OLS PHI premium	(4) OLS PHI premium	(5) OLS Log PHI Premium
Age	4.581***	7.026***	6.267***	6.187***	0.0222***
	(0.661)	(0.299)	(0.286)	(0.306)	(0.00133)
Female	-21.53	129.1***	147.3***	174.7***	0.626***
	(27.00)	(19.94)	(16.96)	(14.98)	(0.0930)
Age x female	-0.0803	-2.100***	-2.469***	-2.993***	-0.0102***
	(0.530)	(0.412)	(0.342)	(0.304)	(0.00181)
Self-employed		117.9***	122.6***	123.0***	0.406***
		(11.67)	(12.24)	(12.76)	(0.0375)
Full-time work		230.4***	157.3***	162.2***	0.477***
		(7.871)	(7.841)	(8.667)	(0.0253)
Single		-31.75***	-30.74***	-30.37***	-0.0788***
		(3.573)	(3.320)	(4.024)	(0.0112)
Log income			67.64***	65.66***	0.195***
			(5.211)	(4.997)	(0.0124)
Diabetes				43.59***	0.106***
				(8.520)	(0.0244)
Asthma				10.88	0.0364
				(9.712)	(0.0287)
Cardiac				35.32**	0.0995***
				(11.33)	(0.0258)
Cancer				1.655	-0.0139
				(13.37)	(0.0342)
Stroke				-2.880	-0.0532
				(26.45)	(0.0825)
Migraine				-18.82*	-0.0349
				(8.052)	(0.0232)
High Blood Pressure				-7.880	-0.0382*
				(6.682)	(0.0171)
Depression				-24.83***	-0.0707*
				(6.754)	(0.0313)
Other diagnosis				19.05***	0.0542**
				(5.072)	(0.0152)
Obese				15.11**	0.0300
				(5.491)	(0.0186)
Disability				-14.91	-0.0383
				(9.397)	(0.0269)
Constant	139.7***	-80.28***	-553.0***	-539.4***	2.886***
	(31.93)	(14.37)	(35.38)	(35.77)	(0.115)
Observations	8502	8488	7579	6264	6311
R^2	0.066	0.299	0.366	0.378	0.402

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.6: Demand for private health insurance

	(1) OLS	(2) Logit MFX	(3) OLS	(4) Logit MFX	(5) OLS	(6) Logit MFX	(7) OLS	(8) Logit MFX
Premium difference PHI-SHI	-0.000832***	-0.000880***	-0.00128***	-0.00148***	-0.000788***	-0.000854***		
	(0.0000510)	(0.0000613)	(0.000147)	(0.000179)	(0.0000594)	(0.0000718)		
Age			0.00191	0.00187			-0.00433**	-0.00477**
			(0.00182)	(0.00202)			(0.00159)	(0.00176)
Female (d)			0.143	0.171			-0.0808	-0.0842
			(0.149)	(0.168)			(0.147)	(0.155)
Age x female			-0.00439	-0.00520			-0.000333	-0.000442
			(0.00318)	(0.00361)			(0.00317)	(0.00353)
Self-employed (d)			0.0233	0.0215			0.245***	0.260***
			(0.0360)	(0.0402)			(0.0247)	(0.0268)
Premium difference PHI-SHI x self-employment			0.000995***	0.00120***				
			(0.000170)	(0.000203)				
Log income			0.0870***	0.0944***			0.109***	0.119***
			(0.0245)	(0.0281)			(0.0221)	(0.0251)
Single (d)			0.0869***	0.0974***			0.108***	0.119***
			(0.0256)	(0.0288)			(0.0256)	(0.0283)
Risk averse in health (d)			-0.0435	-0.0492	-0.0382	-0.0422	-0.0430	-0.0470
			(0.0288)	(0.0327)	(0.0292)	(0.0325)	(0.0296)	(0.0328)
Healthy eating (d)			0.0267	0.0319	0.0191	0.0216	0.0259	0.0287
			(0.0198)	(0.0227)	(0.0198)	(0.0222)	(0.0202)	(0.0225)
Life Insurance (d)			0.0369	0.0434	0.0332	0.0368	0.0361	0.0405
			(0.0205)	(0.0232)	(0.0207)	(0.0229)	(0.0210)	(0.0233)
Employ household help (d)			0.164***	0.188***	0.185***	0.202***	0.166***	0.182***
			(0.0260)	(0.0297)	(0.0257)	(0.0283)	(0.0263)	(0.0288)
Center-left politics (d)			-0.0738**	-0.0885**	-0.0747**	-0.0854**	-0.0767**	-0.0868**
			(0.0277)	(0.0320)	(0.0283)	(0.0319)	(0.0283)	(0.0317)
Higher education (d)			0.0333	0.0381	0.0471*	0.0529*	0.0397	0.0450
			(0.0238)	(0.0270)	(0.0233)	(0.0260)	(0.0241)	(0.0268)
Smoker (d)			0.0354	0.0399	0.0408	0.0454	0.0368	0.0402
			(0.0259)	(0.0295)	(0.0260)	(0.0290)	(0.0263)	(0.0293)
Observations	8129	8129	5701	5701	5701	5701	5701	5701
R^2	0.083		0.133		0.111		0.106	

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

The table provides several specifications for PHI demand. For all specifications the sample includes full-time employees with income above PHI eligibility threshold and self-employed in years 2005-2009. The dependent variable in all specifications is the indicator for PHI enrollment.

Table 3.7: Demand for private health insurance using different income brackets

	(1) OLS Model 1	(2) Logit MFX Model 1	(3) OLS Model 2	(4) Logit MFX Model 2	(5) OLS Model 3	(6) Logit MFX Model 3	(7) OLS Model 4	(8) Logit MFX Model 4
Premium difference PHI-SHI	-0.00128***	-0.00148***	-0.00139***	-0.00154***	-0.00140***	-0.00165***	-0.00238***	-0.00340***
	(0.000147)	(0.000179)	(0.000260)	(0.000288)	(0.000294)	(0.000365)	(0.000267)	(0.000531)
Age	0.00191	0.00187	0.00673*	0.00679*	0.00346	0.00315	0.00821*	0.0116*
	(0.00182)	(0.00202)	(0.00306)	(0.00329)	(0.00347)	(0.00401)	(0.00391)	(0.00545)
Female (d)	0.143	0.171	0.264	0.314	0.212	0.212	0.488	0.568**
	(0.149)	(0.168)	(0.223)	(0.260)	(0.318)	(0.392)	(0.394)	(0.201)
Age x female	-0.00439	-0.00520	-0.00547	-0.00621	-0.00469	-0.00459	-0.0107	-0.0164
	(0.00318)	(0.00361)	(0.00473)	(0.00534)	(0.00686)	(0.00832)	(0.00814)	(0.0108)
Self-employed (d)	0.0233	0.0215	-0.0982	-0.105	0.0685	0.0696	-0.168	-0.289**
	(0.0360)	(0.0402)	(0.0884)	(0.0816)	(0.0821)	(0.0951)	(0.0884)	(0.110)
Premium difference PHI-SHI x self-employment	0.000995***	0.00120***	0.000419	0.000564	0.00128***	0.00156***	0.00187***	0.00274***
	(0.000170)	(0.000203)	(0.000305)	(0.000333)	(0.000336)	(0.000397)	(0.000339)	(0.000553)
Log income	0.0870***	0.0944***	0.336*	0.373*	0.112	0.123	0.430	0.585
	(0.0245)	(0.0281)	(0.155)	(0.173)	(0.226)	(0.266)	(0.351)	(0.460)
Single (d)	0.0869***	0.0974***	0.0776*	0.0875*	0.166***	0.196***	0.0887	0.114
	(0.0256)	(0.0288)	(0.0371)	(0.0417)	(0.0495)	(0.0580)	(0.0597)	(0.0746)
Risk averse in health (d)	-0.0435	-0.0492	-0.0130	-0.0166	-0.0185	-0.0176	-0.0498	-0.0673
	(0.0288)	(0.0327)	(0.0410)	(0.0463)	(0.0504)	(0.0620)	(0.0626)	(0.0812)
Healthy eating (d)	0.0267	0.0319	0.0153	0.0193	0.0442	0.0576	0.0134	0.0251
	(0.0198)	(0.0227)	(0.0276)	(0.0313)	(0.0368)	(0.0446)	(0.0425)	(0.0563)
Life Insurance (d)	0.0369	0.0434	0.0307	0.0371	0.0765*	0.101*	0.111*	0.148*
	(0.0205)	(0.0232)	(0.0303)	(0.0344)	(0.0381)	(0.0446)	(0.0486)	(0.0606)
Employ household help (d)	0.164***	0.188***	0.208***	0.235***	0.136**	0.160**	0.132**	0.167**
	(0.0260)	(0.0297)	(0.0400)	(0.0459)	(0.0468)	(0.0559)	(0.0486)	(0.0609)
Center-left politics (d)	-0.0738**	-0.0885**	-0.0133	-0.0163	-0.0957	-0.116*	-0.0205	-0.0437
	(0.0277)	(0.0320)	(0.0356)	(0.0420)	(0.0494)	(0.0585)	(0.0585)	(0.0797)
Higher education (d)	0.0333	0.0381	0.0437	0.0513	0.0902*	0.106*	0.0809	0.0980
	(0.0238)	(0.0270)	(0.0326)	(0.0368)	(0.0428)	(0.0504)	(0.0519)	(0.0655)
Smoker (d)	0.0354	0.0399	-0.00573	-0.00706	0.0876	0.105	-0.000332	-0.00481
	(0.0259)	(0.0295)	(0.0359)	(0.0409)	(0.0481)	(0.0594)	(0.0587)	(0.0781)
Observations	5701	5701	1950	1950	1236	1236	669	669
R^2	0.133		0.155		0.198		0.255	
Income brackets	all	all	[4000,5000]	[4000,5000]	[5000,6000]	[5000,6000]	[6000,7000]	[6000,7000]

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

The table reports the specifications (3) and (4) from the original demand estimation table conducted on different sub-samples. The dependent variable is the indicator for PHI enrollment. Columns (1) and (2) replicate columns (3) and (4) from the original demand regression.

Table 3.8: Counterfactual consumer surplus changes

	Model 1	Model 2	Model 3	Model 4
Counterfactual 1: open access to PHI				
Change in consumer surplus within sample, mn EUR	0.6	0.2	0.4	0.06
Average change in CS per capita, EUR	39	13	28	4
Change in consumer surplus with sample weights, mn EUR	2,100	730	1,600	240
Counterfactual 2: threshold decreased by 500 EUR				
Percentage of individuals affected	9%	9%	9%	9%
Change in consumer surplus within sample, mn EUR	0.08	0.04	0.06	0.02
Average change in CS per capita, EUR	60	34	48	15
Change in consumer surplus with sample weights, mn EUR	278	155	228	74
Counterfactual 3: threshold decreased by 1000 EUR				
Percentage of individuals affected	24%	24%	24%	24%
Change in consumer surplus within sample, mn EUR	0.2	0.1	0.15	0.04
Average change in CS per capita, EUR	56	28	43	11
Change in consumer surplus with sample weights, mn EUR	737	366	598	157

The table summarizes calculations for the counterfactual changes in consumer welfare from changes in the enrollment eligibility *ceteris paribus*. Four columns refer to four demand models that were estimated using the same specification on different income subsamples. (See demand estimation tables). The key differences across the models comes from differential predictions of how many individuals currently below the eligibility threshold would be willing to switch to the PHI if they were given the opportunity, which in turn come from different representative utility predictions. In each counterfactual calculation the first row reports the fraction of individuals that would want to switch as a percent of all individuals that would have become eligible. The switchers are calculated by counting all individuals that have positive predicted utility from the PHI. The second row divides the total sum of the predicted positive utilities by the number of switchers. The third row weights all positive utilities by the sample weights for the respective individuals and adds these up. All calculations use observations only on full-time employees with income below the current PHI eligibility threshold.

Table 3.9: Mean values of observed covariates for full-time employees

Variable	All full-time employees	PHI-eligible	PHI-eligible with SHI	PHI-eligible with SHI	t-statistic PHI/SHI
Income	3036	5789	5595	6141	-7.79
PHI	.08	.35	0	1	-
Age	43	46	47	46	3.48
Health status (1=excellent, 5=bad)	2.5	2.4	2.4	2.27	5.3
Female	.33	.15	.15	.17	-2.7
Number outpatient visits	1.8	1.7	1.7	1.6	1.8
Probability outpatient visit	.62	.61	.62	.58	2.6
Number inpatient visits	.09	.09	.1	.07	2.2
Probability inpatient visits	.07	.07	.07	.07	0.6
Older than 50	.26	.33	.34	.304	3.02
Older than 60	.03	.04	.047	.046	0.3
Younger than 35	.23	.07	.07	.07	-0.94
BMI	26	26	26.4	25.8	3.4
Smoker	.34	.22	.22	.22	-0.25
Risk aversion in health (0=risk averse)	3.06	3.47	3.44	3.52	-1.09
Risk aversion (0=risk averse)	4.7	5.1	5.05	5.28	-3.4
Disability	.06	.05	.0668	0.03	5.7
Healthy eating (1=always, 4=never)	2.6	2.5	2.6	2.46	4.93
Sport (1=never, 4=often)	2.5	2.9	2.9	3.09	-5.51
Health satisfaction (10=most satisfaction)	6.9	7.1	7.02	7.25	-4.22

The table provides basic summary statistics for the sample of all full-time employees in 2005-2009. Columns (1)-(4) report means of the respective covariates. Column (1) includes all observations. Column (2) includes all employees with income above the eligibility threshold. Columns (3) and (4) report the means of the latter groups differentially for those employees who chose PHI and those who chose SHI. Column (5) provides a two-tailed mean comparison t-test for the last two groups.

Table 3.10: Covariate balance: means of observed covariates for full-time employees within 500EUR bandwidth around the threshold

Variable	500 EUR below threshold	500 EUR above threshold	t-test difference in means
Income	3697	4182	-73.9
Age	44.2	44.4	-0.6
Health status (1=excellent, 5=bad)	2.4	2.4	0.9
Female	0.25	0.21	2.5
Number outpatient visits	1.8	1.8	0.4
Probability outpatient visit	0.63	0.62	0.4
Number inpatient visits	0.09	0.09	-0.3
Probability inpatient visits	0.06	0.06	0.3
Older than 50	0.265	0.255	0.6
Older than 60	0.024	0.025	-0.06
Younger than 35	0.15	0.13	1.8
BMI	26.1	26.2	-0.38
Smoker	0.31	0.24	3.7
Risk aversion in health (0=risk averse)	3.2	3.3	-1.5
Risk aversion (0=risk averse)	4.8	4.8	-0.45
Disability	0.06	0.07	-1.1
Healthy eating (1=always, 4=never)	2.6	2.6	-0.844
Sport (1=never, 4=often)	2.65	2.7	-0.95
Health satisfaction (10=most satisfaction)	6.9	7.0	-0.83

The table provides a basic covariate balance summary statistics for the sample of all full-time employees in 2005-2009. Column (1) reports the means of the respective covariates for all full-time employees with income level within a 500 EUR window below the cutoff. Column (2) reports the same for the employees with income within a 500 EUR window above the cutoff. Column (3) provides a two-tailed mean comparison t-test for the two groups. The comparison bandwidth is relatively wide due to the sparsity of data at smaller bandwidth restrictions.

Figure 3-1: Local-linear smoothing of several diagnosis probabilities by insurance type for all income levels

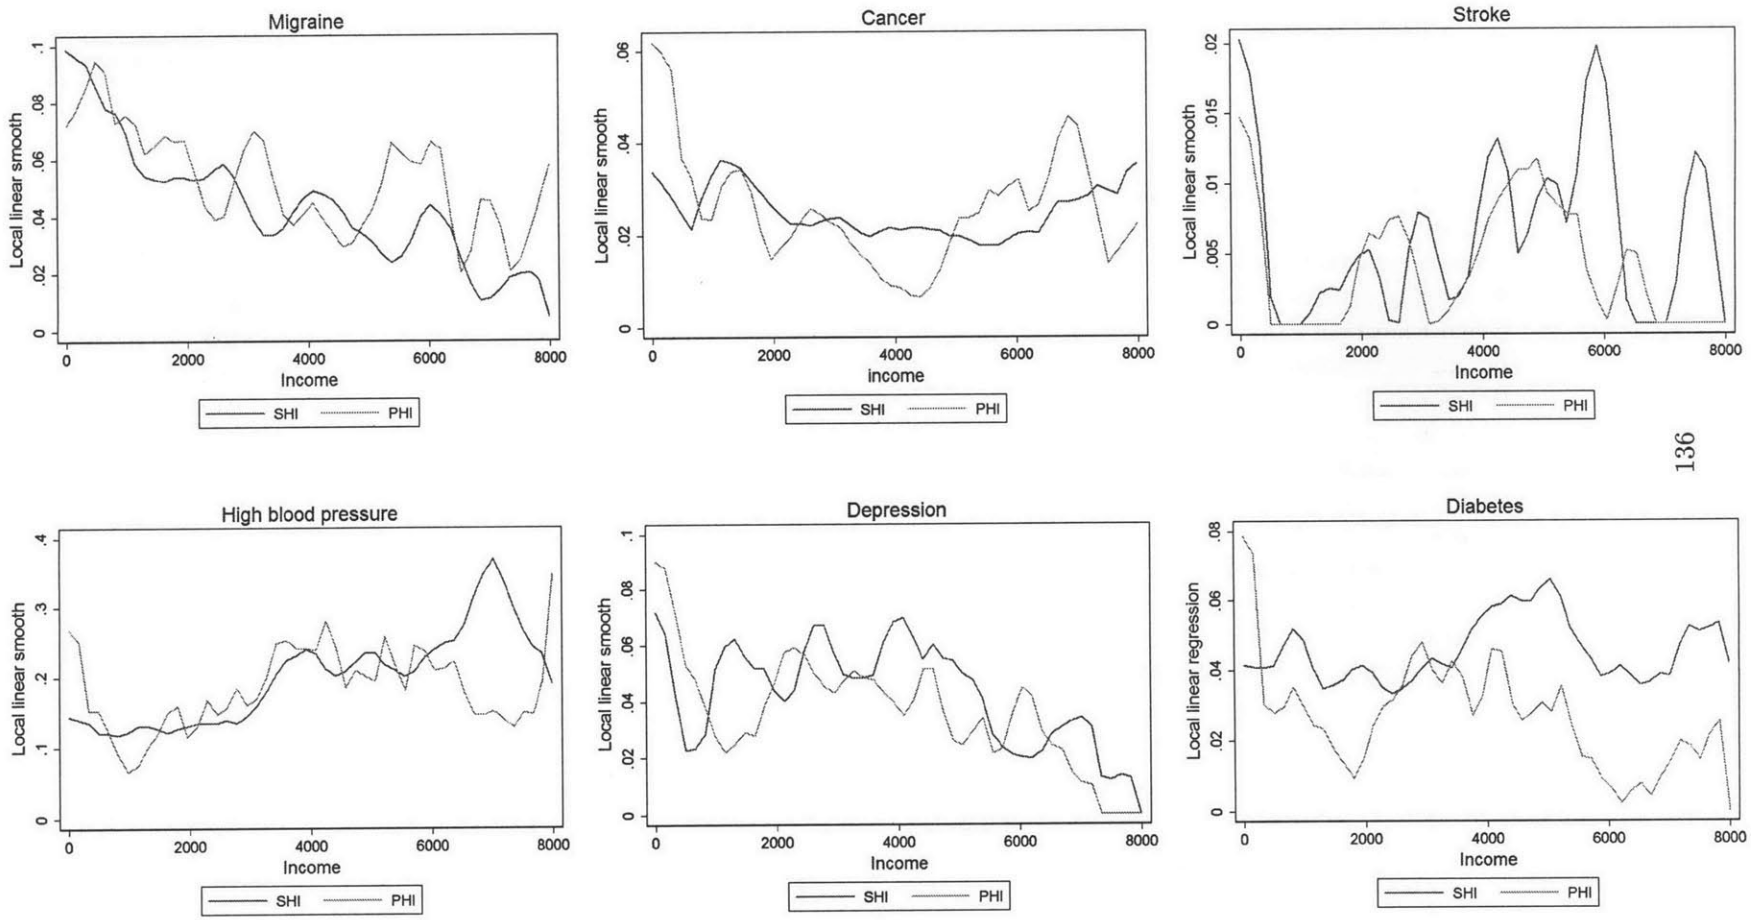

Figure 3-2: Characteristics of the SHI-insured above and below PHI-eligibility threshold

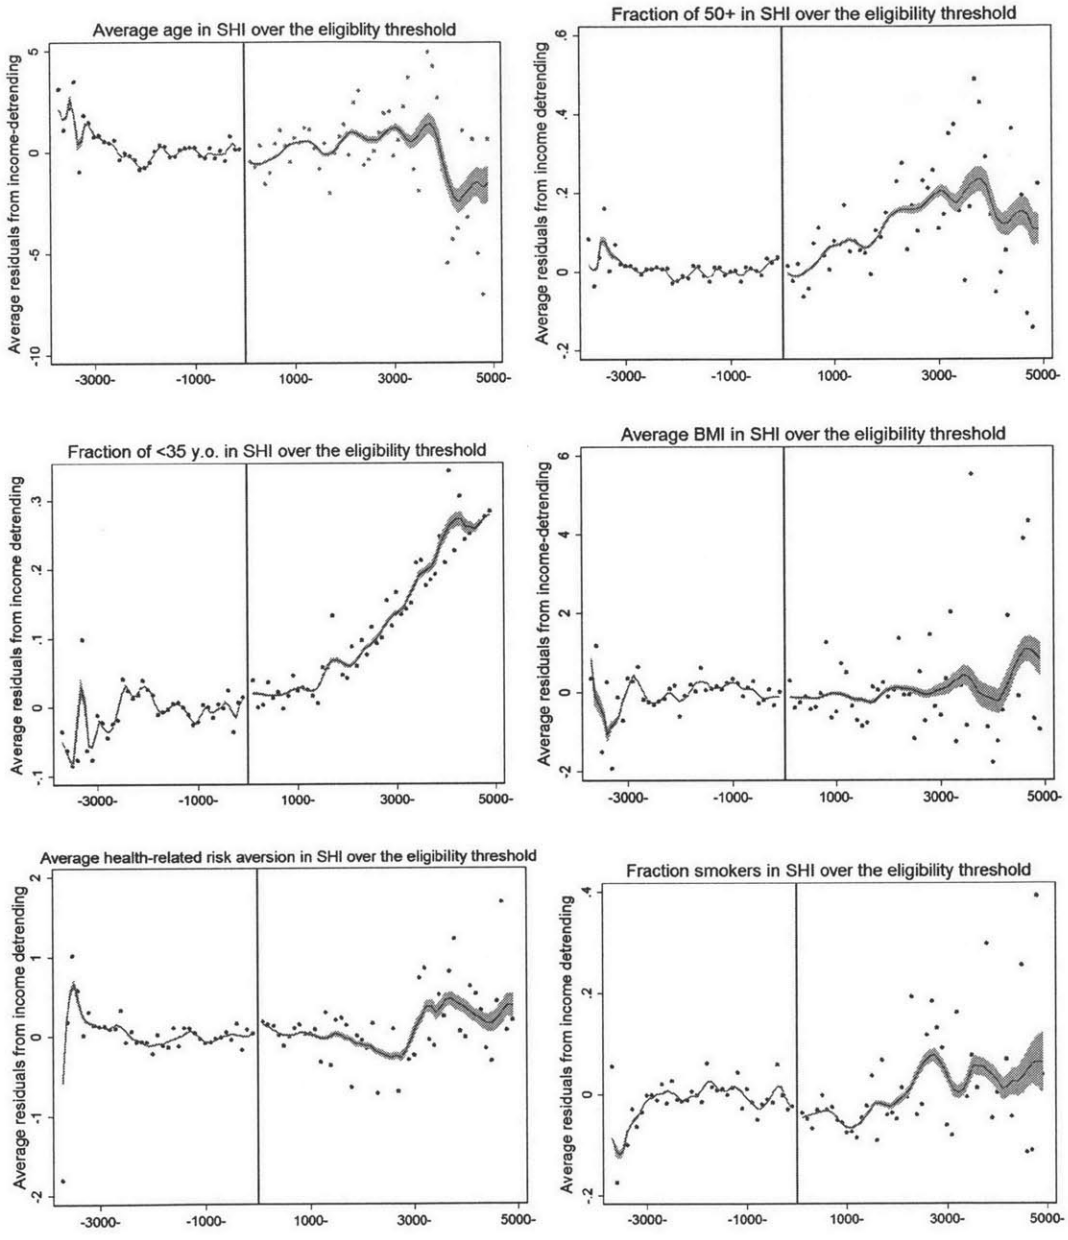

The panels aim to graphically detect the presence or lack of a significant discontinuity in the characteristics of the SHI-insured at the crossing of the PHI-eligibility threshold. Since many of the characteristics are correlated with income, the covariates were first regressed on the income categories in order to eliminate the natural income-induced trend. The regressions did not allow for breaks at the income eligibility threshold. Then, average residuals per income bin were computed and plotted separately to the left and to the right of the cutoff. The sample includes all employees in years 2005-2009 that reported SHI enrollment.

Figure 3-3: Characteristics of the SHI-insured above and below PHI-eligibility threshold, continued

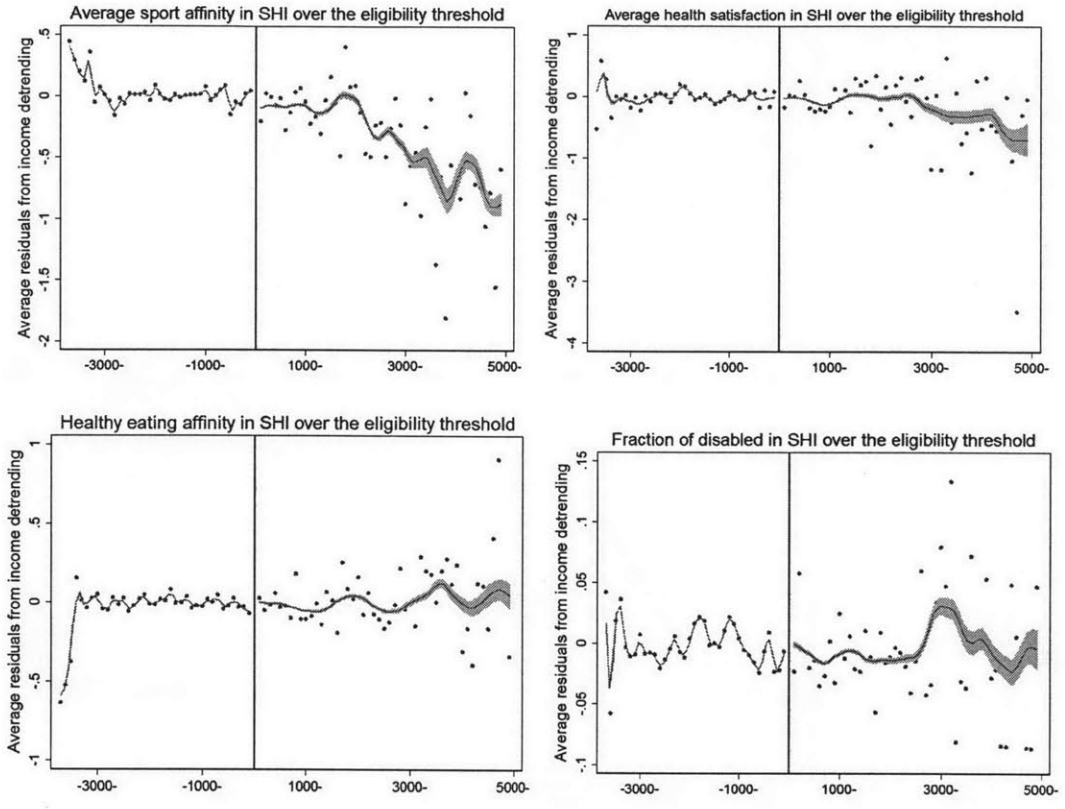

The panels aim to graphically detect the presence or lack of a significant discontinuity in the characteristics of the SHI-insured at the crossing of the PHI-eligibility threshold. Since many of the characteristics are correlated with income, the covariates were first regressed on the income categories in order to eliminate the natural income-induced trend. The regressions did not allow for breaks at the income eligibility threshold. Then, average residuals per income bin were computed and plotted separately to the left and to the right of the cutoff. The sample includes all employees in years 2005-2009 that reported SHI enrollment.

Figure 3-4: First stage: graphical representation of the relationship between PHI-eligibility and PHI enrollment

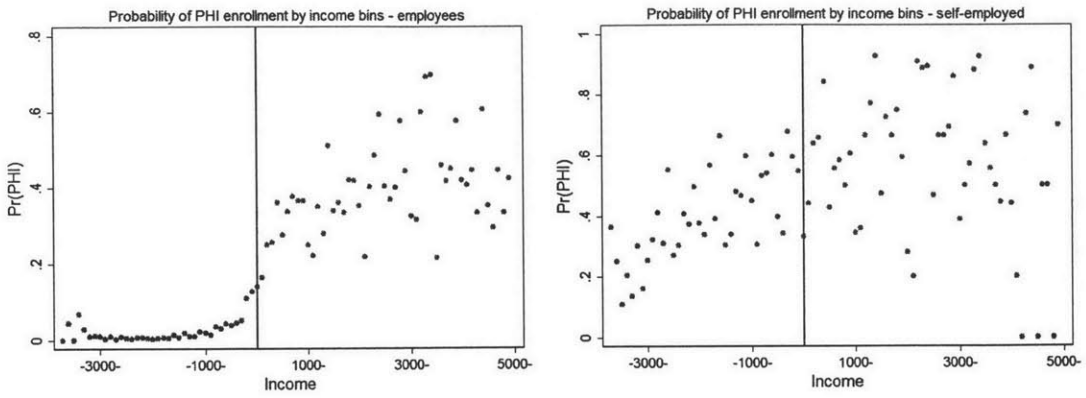

The left panel shows the employees' probability of PHI-enrollment by equally-sized income bins of 100 EUR. Income refers to self-reported monthly pre-tax employment income. All full-time employees of the 2005-2009 survey years are included. Income is centered around the eligibility threshold for the respective years. The reference line indicates the eligibility threshold. The right panel shoes the same PHI-enrollment probability by income bins for the sub-sample of the self-employed in years 2005-2009. The PHI eligibility threshold rule does not apply to the self-employed.

Figure 3-5: Reduced form: graphical representation of the relationship between healthcare utilization and PHI-eligibility

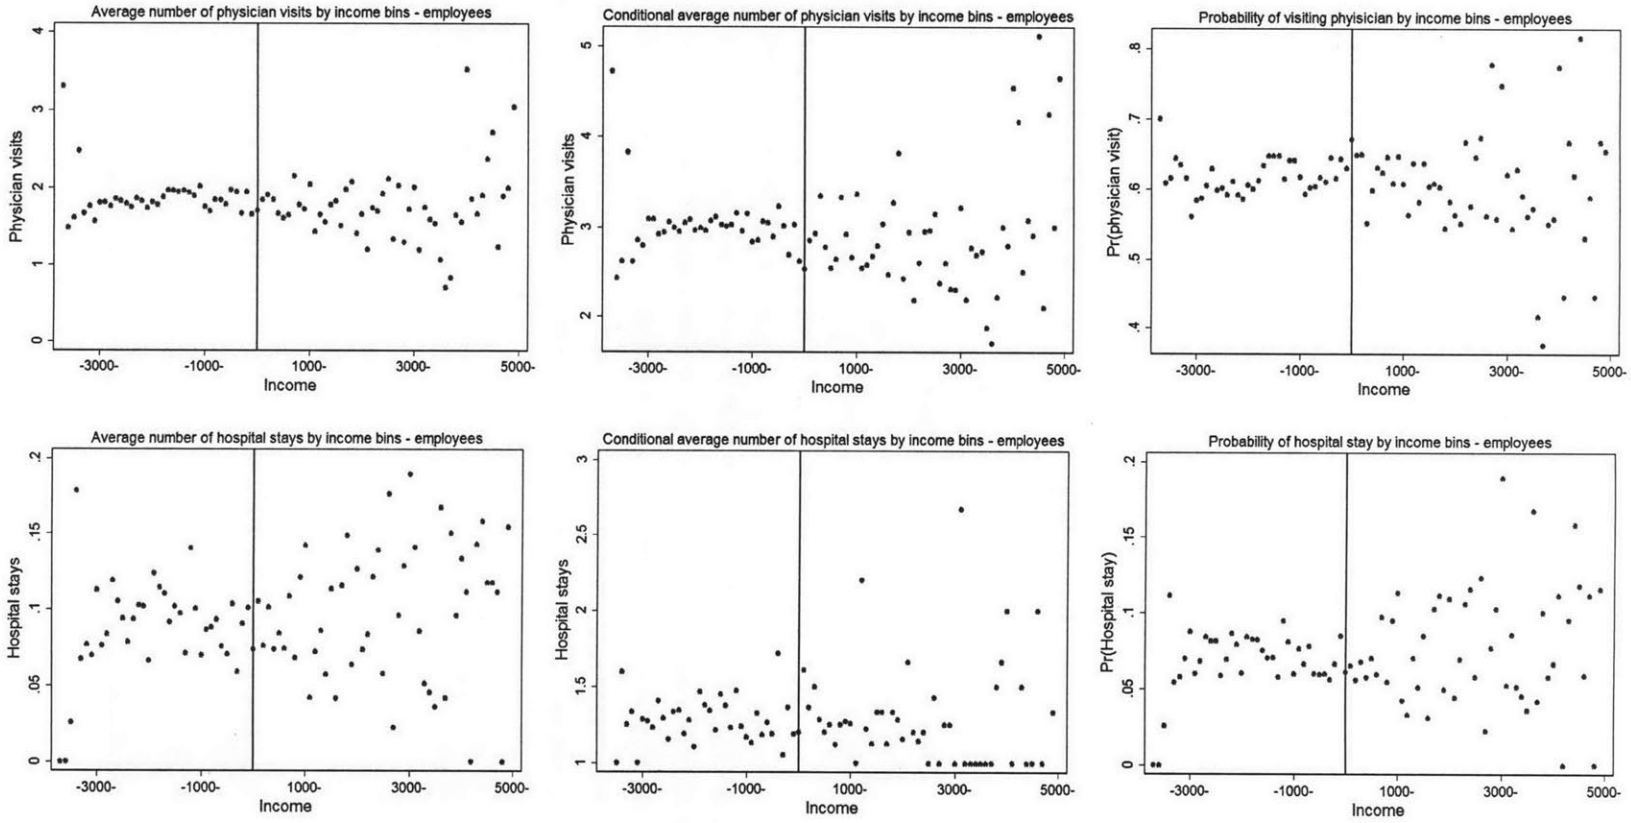

The sample includes all full-time employees in 2005-2009. The income is centered at the PHI eligibility threshold and divided into equally spaced bins of 100 EUR. For each bin, the average of several measures of healthcare utilization is plotted. The reference line corresponds to the PHI eligibility threshold. The threshold was computed individually. For most individuals it just corresponds to the official eligibility threshold in years 2005-2009 respectively. For individuals that are observed to have had PHI before 2005, a lower threshold is applied as was specified by the regulation after the changes in the threshold in 2005. The healthcare utilization measures are the same as the ones used in the reduced form regressions.

Figure 3-6: In-sample and out-of-sample predictions of the linear PHI pricing approximation

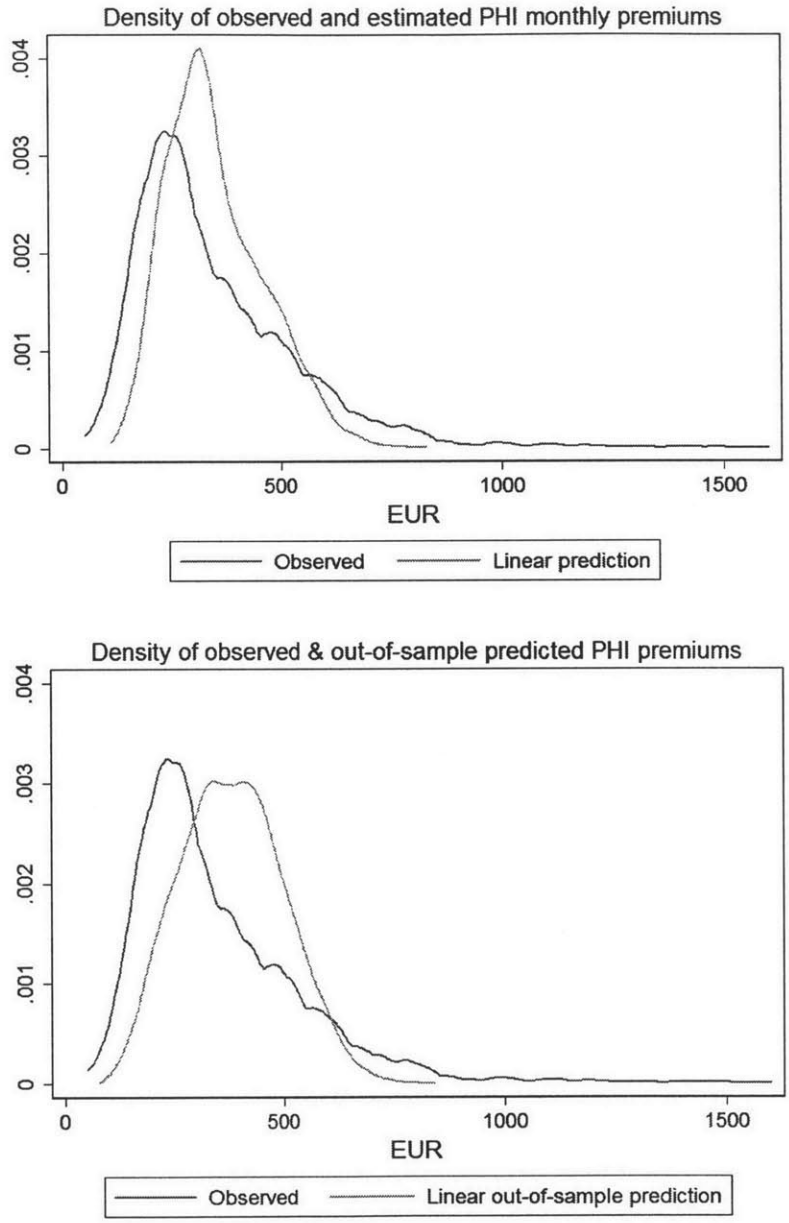

The four panels present different kernel density estimates of the goodness-of-fit and out-of-sample predictions calculated using model (5) of the linear PHI pricing specification. Left-top panel compares the observed premiums and the calculated in-sample predictions. The sample includes all individuals in 2005-2009 that reported PHI monthly premiums. Right-top panel compares the distributions of the observed premiums and the premium predictions for all PHI-eligible but SHI-insured individuals. The bottom two panel provide the same comparisons, but the sample is restricted to full-time employees only.

Figure 3-7: Distribution of observed and imputed PHI premiums

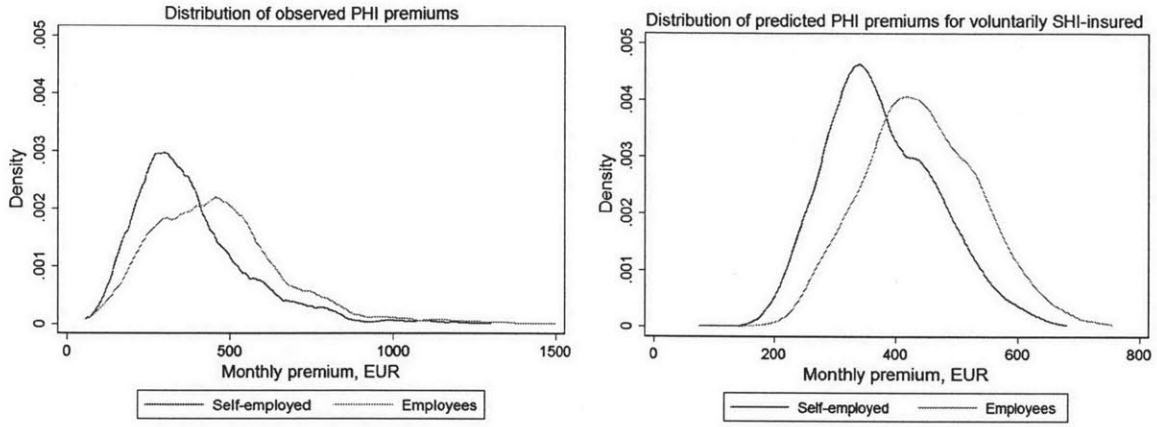

The left panel shows a kernel estimation of the empirical density for the PHI premiums reported in the data. The density is shown separately for the self-employed and the employees. Note that a critical assumption about the data here is that the premium reported by the employees are the actual payments by the employees and do not include the employer subsidy. Unfortunately, the way the question was formulated in the survey does not allow any check on whether this is the right assumption. The right panel shows kernel estimates of densities for premiums that were imputed for the SHI-insured individuals that could have chosen to opt out into the PHI. The predictions followed model (5) of the linear PHI pricing specification.

Figure 3-8: SHI premiums and implied PHI-SHI premium differences

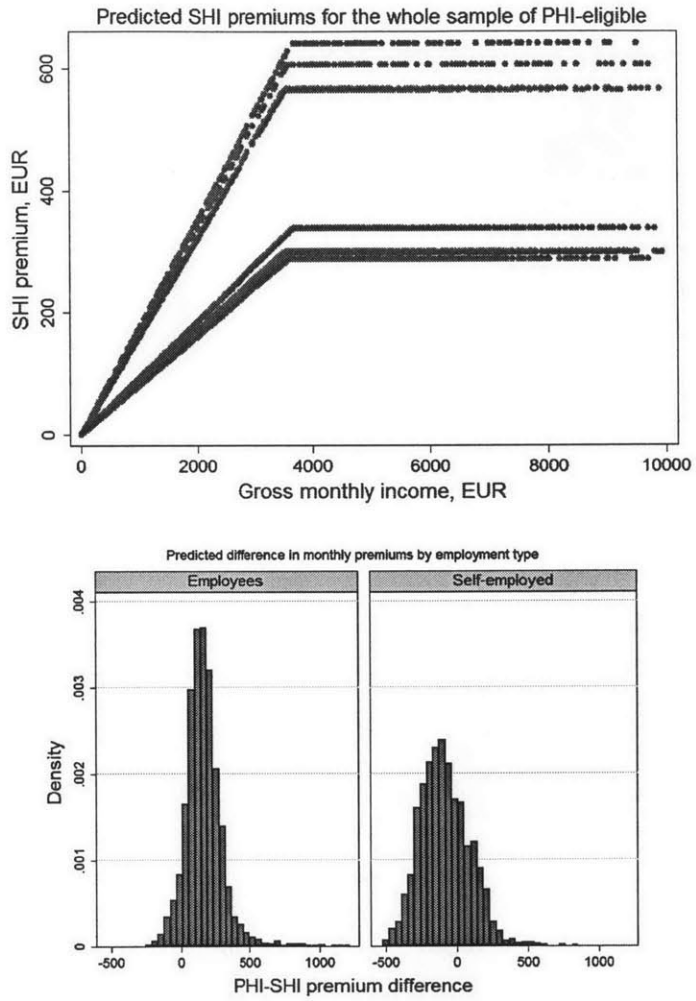

The top-left panel shows the individual predictions of SHI premiums for all PHI-eligible full-time employees and the self-employed. The rule for the SHI premium calculations relies on the SHI regulations for 2005-2009. The SHI premium includes long-term-care insurance premiums, so as to facilitate the comparison to the PHI premiums that usually include the long-term-care tariff. For 2005, the SHI premium was calculated as 8.45% of income for employees and 16% for the self-employed, with caps at 297.86 and 564 EUR respectively. For 2006 and 2007, the percentages are 8.4% and 15.9% with caps at 299.25 and 566.43 EUR. For 2008, the percentages are 7.975% and 16.85%, capped at 287.1 and 606.6 EUR. Finally, for 2009 the rates are 9.175% and 17.45% respectively, capped at 337.18 EUR and 641.26 EUR. The top-right panel plots the imputed differences that PHI-eligible employees faced for a PHI vs. SHI coverage. The SHI premiums were computed as above, while PHI premiums were either taken as reported or imputed from the pricing function (model 5). The densities are reported separately by the observed choice of insurance. The bottom panel splits the predicted differences in the premiums by the employment type rather than by the choice of insurance.

Figure 3-9: Distribution of predicted PHI representative utility relative to the SHI for PHI-eligible individuals

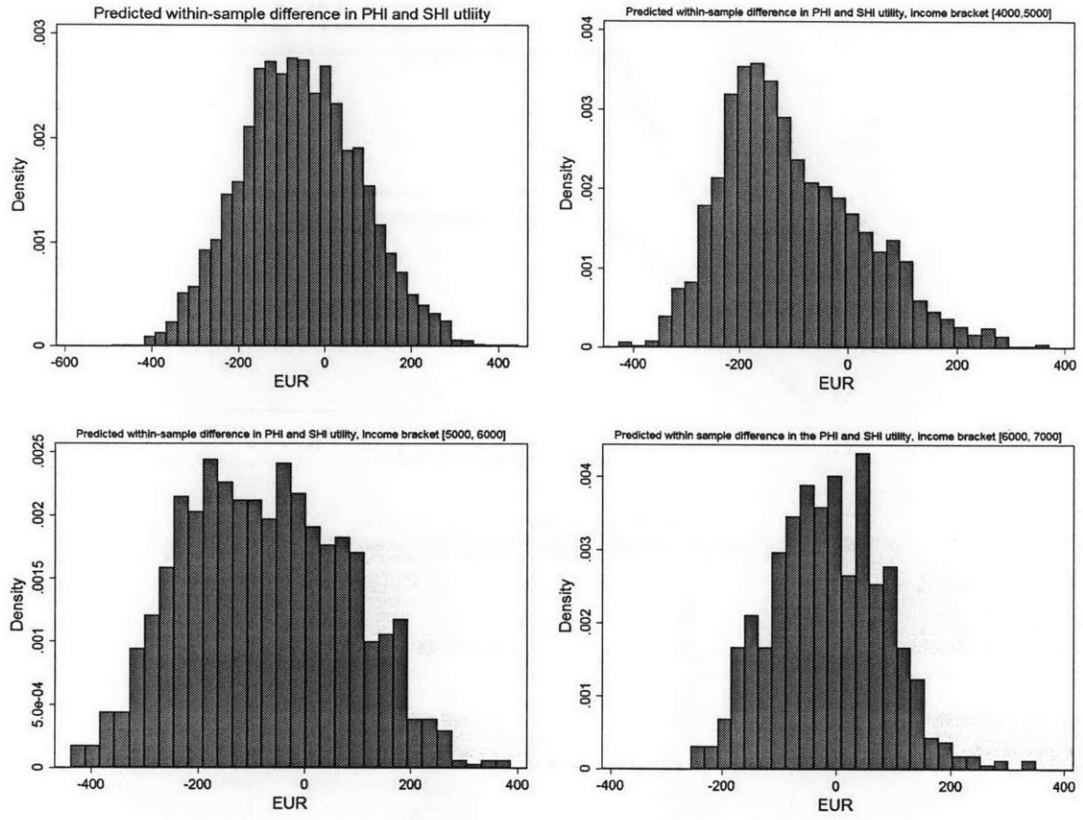

The panels show the distribution of the difference in the predicted levels of the representative utility ($V_{iPHI} - V_{iSHI}$). The sample includes all PHI eligible employees and the self-employed in years 2005-2009. Four panels correspond to the four different demand models estimated on four income sub-samples. That is, the top-left panel was constructed using the estimated coefficients from the logit specification on the full sample of PHI-eligible employees and self-employed of all income levels. The top-right panel uses coefficients from the logit specification on the individuals with income between 4000 and 500 EUR. Since there are only two alternatives, the difference in the representative utility estimate is constructed by simply multiplying the matrix of the logit coefficients with the matrix of individual characteristics. To represent the utility in monetary terms, the demand coefficients were first multiplied by the inverse of the coefficient on price. The representation does not account for the unobserved stochastic component of utility.

Figure 3-10: Estimated gain in expected consumer surplus for consumers of different income levels from the policy change

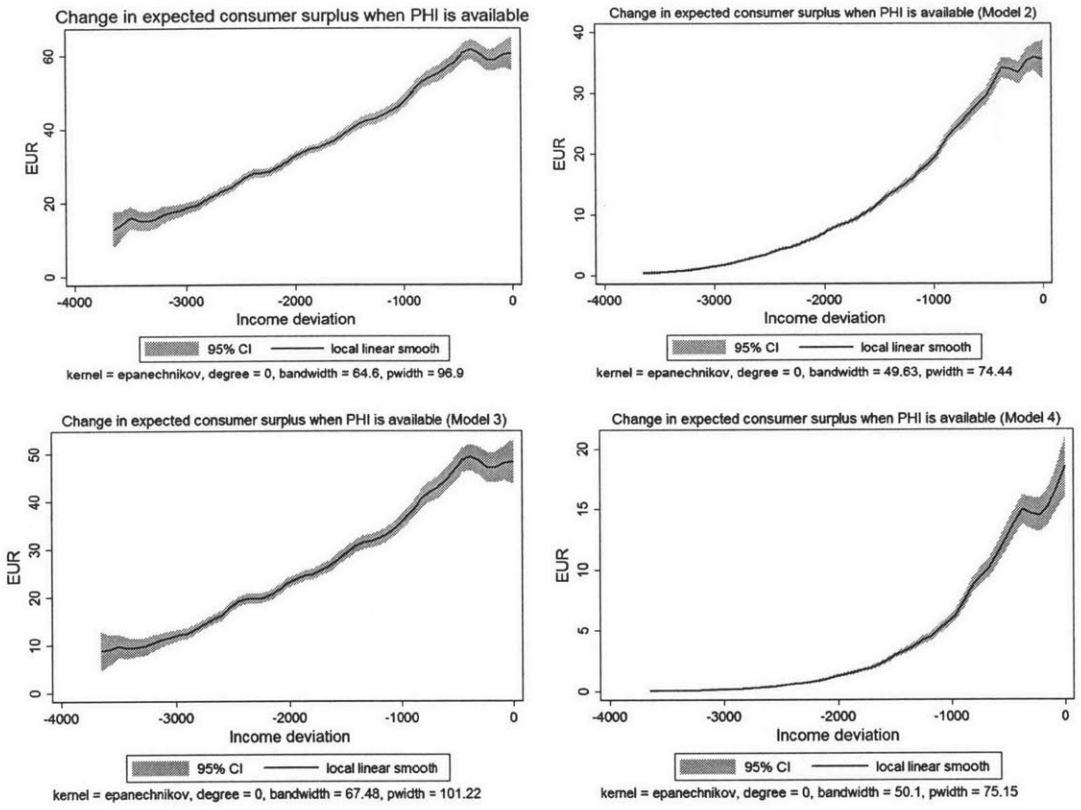

The graphs represent the average per caputa changes in expected consumer surplus at different income levels. The changes in expected surplus were calculated for each individual using $\triangle E(CS_i) = \frac{1}{\hat{\gamma}}[ln(1+e^{\hat{V}_{iPHI-SHI}})]$. To show the difference in the consumer surplus changes across different income categories, the consumer surplus values were plotted against the deviations from the income threshold using local linear regression. The sample includes employees with monthly income below the PHI eligibility threshold.

Figure 3-11: Changes in expected consumer surplus for different policy counterfactuals

1. Changes in the expected consumer sample within the observed sample; individual observations were not weighted

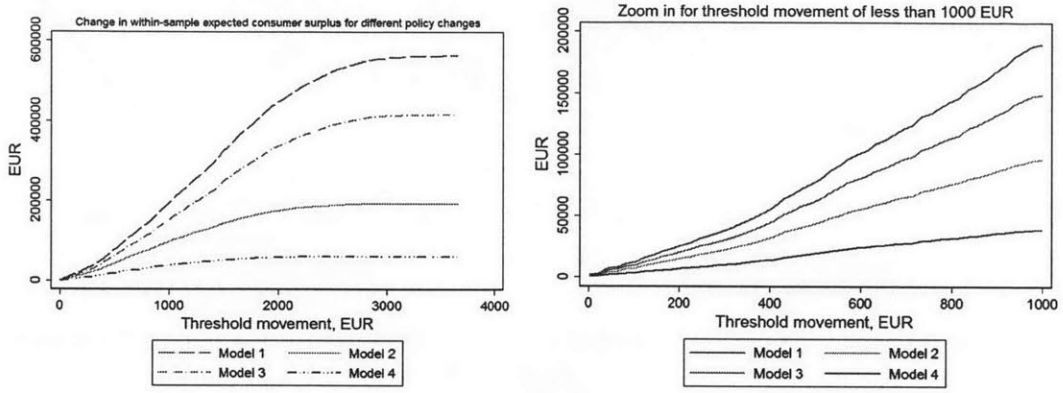

2. Changes in the expected consumer sample for the population; individual observations were weighted with the sample weights of the survey

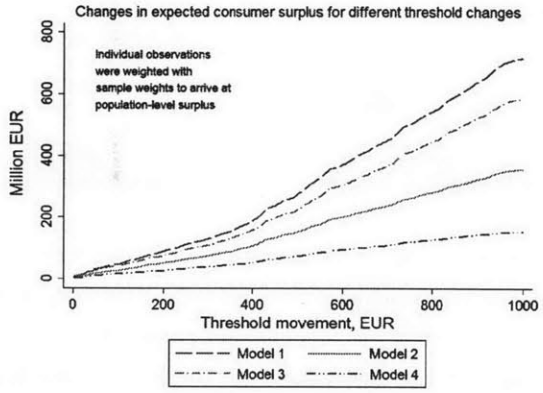

These graphs were constructed by adding up changes in the individual consumer surplus as calculated with $\triangle E(CS_i) = \frac{1}{\hat{\gamma}}[ln(1 + e^{\hat{V}_{iPHI-SHI}})]$ for different levels of changes in the income threshold policy. That is, to get the welfare effect of moving the threshold by 200 EUR, changes in the expected consumer surplus of the individuals with income that is within 200 EUR of the threshold was added up. The last graph adds up individual consumer surplus calculations weighted by population sample weights of the survey.

Bibliography

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

Abaluck, J., Gruber, J., 2013. Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance. NBER Working Paper no. 19163.

Angrist, J., Pischke, J.-S., 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Armstrong, M., Sappington, D. E. M., 2007. Recent Developments in the Theory of Regulation. In: Armstrong, M., Porter, R. (Eds.), Handbook of Industrial Organization, Vol. 3. No. 06. Elsevier B.V., pp. 1557–1700.

Bauhoff, S., 2012. Do Health Plans Risk-Select? An Audit Study on Germany's Social Health Insurance. Journal of Public Economics v. 96, 750–759.

Berry, S., Levinsohn, J., Pakes, A., 1995. Automobile Prices in Market Equilibrium. Econometrica 63(4), 841–889.

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., 2013. Who Uses the Roth 401(k), and How Do They Use It? NBER Working Paper No. 19193.

Brown, J., Duggan, M., Kuziemko, I., Woolston, W., 2012. How does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program. Unpublished mimeo.

Bundorf, K., Levin, J., Mahoney, N., 2012. Pricing and Welfare in Health Plan Choice. American Economic Review 102(7), 1–38.

Chetty, R., Friedman, J. N., Leth-Petersen, S., Nielsen, T., Olsen, T., 2013. Active vs. Passive Decisions and Crowdout in Retirement Savings Accounts: Evidence from Denmark. NBER Working Paper No. 18565.

Chiappori, P.-A., Salanie, B., 2000. Testing for Asymmetric Information in Insurance Markets. The Journal of Political Economy 108 (1), 56–78.

Cutler, D., Finkelstein, A., McGarry, K., 2008. Preference Heterogenity and Insurance Markets: Explaining a Puzzle of Insurance. American Economic Review Papers and Proceedings 98(2), 157–162.

Decarolis, F., 2013. Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design? Mimeo, Boston University.

Dube, J.-P., Hitsch, G. J., Rossi, P. E., 2010. State Dependence and Alternative Explanations for Consumer Inertia. RAND Journal of Economics 41, 417–445.

Duggan, M., Healy, P., Scott Morton, F., 2008. Providing Prescription Drug Coverage to the Elderly: America's Experiment with Medicare Part D. Journal of Economic Perspectives 22(4), 69–92.

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

Einav, L., Finkelstein, A., Levin, J., 2010a. Beyond Testing: Empirical Models of Insurance Markets. Annual Review of Economics 2, 311–336.

Einav, L., Finkelstein, A., Schrimpf, P., 2010b. Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence From the U.K. Annuity Market. Econometrica 78(3), 1031–1092.

Einav, L., Finkelstein, A., Schrimpf, P., 2013. The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D. NBER Working Paper no. 19393.

Ericson, K., 2013. Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange. American Economic Journal: Economic Policy (forthcoming).

Farrell, J., Klemperer, P., 2007. Coordination and Lock-In: Competition with Switching Costs and Network Effects. In: Armstrong, M., Porter, R. (Eds.), Handbook of Industrial Organization, Vol. 3. No. 06. Elsevier B.V., pp. 1967–2072.

Finkelstein, A., 2004. Minimum standards, insurance regulation and adverse selection: evidence from the Medigap market. Journal of Public Economics 88 (12), 2515–2547.

Finkelstein, A., McGarry, K., 2006. Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market. American Economic Review 96(4), 938–958.

Geruso, M., 2013. Selection in Employer Health Plans: Homogeneous Prices and Heterogeneous Preferences. Mimeo, Harvard University and UT Austin.

Grunow, M., Nuscheler, R., 2013. Public and Private Health Insurance in Germany: the Ignored Risk Selection Problem. Health Economics May 21.

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

Heckman, J., 1981. Structural Analysis of Discrete Data with Econometric Applications. Cambridge, Mass.: MIT Press, Ch. Statistical Models for Discrete Panel Data, pp. 114–78.

Heckman, J., 1991. Identifying the Hand of Past: Distinguishing State Dependence from Heterogeneity. American Economic Review 81(2), 75–79.

Heckman, J., Singer, B., 1986. Handbook of Econometrics. Handbooks in Economics series, book 2 Amsterdam; Oxford and Tokyo: North-Holland;, Ch. Econometric Analysis of Longitudinal Data, pp. 1689–763.

Heiss, F., Leive, A., McFadden, D., Winter, J., 2013. Plan Selection in Medicare Part D: Evidence from Administrative Data. Journal of Health Economics (forthcoming).

Heiss, F., McFadden, D., Winter, J., 2009. Regulation of Private Health Insurance Markets: Lessons from Enrollment, Plan Type Choice, and Adverse Selection in Medicare Part D. NBER Working Paper No. 15392.

Heiss, F., Mcfadden, D., Winter, J., 2010. Mind the Gap ! Consumer Perceptions and Choices of Medicare Part D Prescription. In: Wise, D. (Ed.), Research Findings in the Economics of Aging. The University of Chicago Press, pp. 413–481.

Ho, K., Hogan, J., Scott Morton, F., 2013. Insurer (In)Efficiency in the Medicare Part D Program. Mimeo, Columbia University and Yale University.

Hoadley, J., Cubanski, J., Hargrave, E., Summer, L., Huang, J., 2012. Medicare Policy data spotlight Medicare Part D: A First Look at Part D Plan Offerings in 2013. KFF Report (October 2012).

Hole, A. R., 2007. Fitting Mixed Logit Models by Using Maximum Simulated Likelihood. The Stata Journal 7-3, 388–401.

Honoré, B., 2002. Nonlinear Models with Panel Data. Portuguese Economic Journal 1(2), 163–179.

Honoré, B., Tamer, E., 2006. Bounds on Parameters in Panel Dynamic Discrete Choice Models. Econometrica 74(3), 611–629.

Hullegie, P., Klein, T., 2010. The Effect of Private Health Insurance on Medical Care Utilization and Self-Assessed Health in Germany. Health Economics Sep;19(9), 1048–62.

Joskow, P., Rose, N., 1989. The Effects of Economic Regulation. In: Schmalensee, R., Willig, R. (Eds.), Handbook of Industrial Organization, Vol. II. Elsevier Science Publishers B.V., pp. 1449–1506.

Kesternich, I., Heiss, F., McFadden, D., Winter, J., 2013. Suit the Action to the Word, the Word to the Action: Hypothetical Choices and Real Decisions in Medicare Part D. Journal of Health Economics (2013).

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

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

Kuziemko, I., Meckel, K., Rossin-Slater, M., 2013. Do Insurers Risk-Select Against Each Other? Evidence from Medicaid and Implications for Health Reform. NBER Working Paper No. 19198.

Lucarelli, C., Prince, J., Simon, K., 2012. The Welfare Impact of Reducing Choice in Medicare Part D: A Comparison of Two Regulation Strategies. International Economic Review 53 (4), 1155–1177.

Lustig, J., 2011. Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare. Unpublished mimeo.

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

Nosal, K., 2012. Estimating Switching Costs for Medicare Advantage Plans. Mimeo, University of Mannheim.

Nuscheler, R., Knaus, T., 2005. Risk Selection in the German Public Health Insurance System. Health Economics 14(12), 1253–71.

Petrin, A., Train, K., 2009. A Control Function Approach to Endogeneity in Consumer Choice Models. Journal of Marketing Research XLVI.

Schmitz, H., 2011. Direct Evidence of Risk Aversion as a Source of Advantageous Selection in Health Insurance. Economic Letters 113(2), 180–182.

Train, K., 2003. Discrete Choice Methods with Simulation. Cambridge University Press.

Train, K., 2009. Discrete Choice Methods with Simulation, 2nd Edition. Cambridge University Press.

Villas-Boas, J. M., Winer, R., 1999. Endogeneity in Brand Choice Models. Management Science 45(10), 1324–1338.

Yin, W., Basu, A., Zhang, J. X., Rabbani, A., Meltzer, D. O., Alexander, G. C., 2008. The Effect of the Medicare Part D Prescription Benefit on Drug Utilization and Expenditures. Annals of Internal Medicine 148, 169–177.

Yin, W., Lakdawalla, D., 2010. Insurers' Negotiating Leverage and the External Effect of Medicare Part D. NBER Working Paper No. 16251.