Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts

Benjamin Handel[*]

January 26, 2010

Abstract

This paper investigates consumer switching costs in the context of health insurance markets, where adverse selection is a potential concern. Switching costs contribute to poor choices when the market environment changes and consumers do not adjust appropriately. Though previous work has studied the problems of adverse selection and consumer choice inadequacy in isolation, these phenomena interact in a way that directly impacts market outcomes. We use a unique proprietary panel data set with the health plan choices and medical utilization of employees at a large firm to show that (i) switching costs are large and (ii) switching costs significantly impact the degree of adverse selection in equilibrium. We estimate a structural choice model to jointly quantify switching costs, risk preferences, and health risk in the population. We use the output of this model to study the welfare impact of an information provision policy that nudges consumers toward better decisions by reducing switching costs. In a partial equilibrium setting where observed plan prices are held fixed, we find that a policy that completely eliminates switching costs improves consumer welfare by 10%. In a general equilibrium setting where insurers change prices to reflect the expenses of their risk pools, the same policy (i) exacerbates adverse selection (ii) *reduces* consumer welfare by 6% and (iii) has distributional implications that favor those who switch as a result of the intervention relative to those who do not.

JEL CODES: D81, D82, D83, G22, I11, I18

Keywords: Adverse selection, switching costs, health insurance, information provision

1 Introduction

The market for health insurance in the United States covers approximately $2 trillion dollars in medical expenditures each year, about 15% of national GDP. Despite the evident value consumers place on such insurance, many individuals are unable to secure coverage at rates that make insurance purchase worthwhile. Indeed, the fact that over 20% of Americans under 65 are uninsured has led to the current push for national health insurance reform.

A number of potential impediments stand in the way of efficiency in the health insurance market. The most noted of these is adverse selection. In competitive markets, prices reflect the expected risk (costs) of the insured pool. Absent the ability to price all risk characteristics, riskier consumer choose more comprehensive health plans, causing the equilibrium prices of these plans to rise and healthier enrollees to select less comprehensive coverage. Previous work has found mixed empirical evidence on the existence and welfare consequences of adverse selection in health insurance markets. Carlin and Town (2007) find evidence of adverse selection but argue that it has minimal welfare consequences, while Einav, Finkelstein, and Cullen (2009) show evidence that adverse selection leads to welfare losses but standard policy instruments may not be able to rectify the problem in a cost-effective manner. Citing evidence from five different markets with the potential for adverse selection, Cutler, Finkelstein, and McGarry (2008) reveal that selection occurs on multiple dimensions, such as preference heterogeneity, and, as a result, may not always be adverse.

A second, but much less studied, potential impediment is poor decision making by consumers. A collection of research summarized by Sunstein and Thaler (2008) presents strong evidence that consumer decisions are heavily influenced by context and can systematically depart from those that would be made in a rational frictionless environment. Madrian and Shea (2001) reveal that the 401(k) choices of consumers are strongly influenced by the default option, leading to starkly different outcomes across different choice environments where a lot of money is at stake. These decision making issues may be magnified when the costs and benefits of each option are difficult to evaluate, as in the market for health insurance.[1]

In fact, in health insurance markets, these two problems can interact in significant ways, because choice can directly affect the extent of adverse selection. As a result, policies designed to improve consumer choice may have an ambiguous welfare effect as the impact of better decision making conditional on prices is traded off with potentially increased adverse selection. This stands in stark contrast to previous work on choice inadequacy where policies designed to improve consumer choices can only have a positive welfare impact.[2]

In this paper we explore the presence of, and interactions between, the problems of adverse selection and consumer choice inadequacy using a unique panel dataset that contains the health

---

[1]In the context of the recent health reform debate, advocates of health insurance exchanges cite policies to help consumers make better decisions as a key component of market regulation (see e.g. Enthoven, Garber, and Singer (2001)).

[2]One instance of previous work of this kind in health insurance is Abaluck and Gruber (2009) who study the choice adequacy of seniors for Medicare Part D prescription drug plans but do so in a partial equilibrium environment where plan prices are held fixed.

insurance choices and subsequent claims realizations for all employees at a large Midwestern firm. Specifically, we study consumer inertia or switching costs, which is one potential source of choice inadequacy. The structure of plan offerings at the firm coupled with the detailed utilization information provides an unusual opportunity to investigate switching costs and adverse selection in tandem. There are two unique elements of the plan choice structure that allow us to go beyond previous work to investigate these phenomena. First, in the middle of the time period studied the firm significantly changed the menu of health plan options. At the time of the menu change, the firm forced all consumers out of their old plan options and made them select a plan from the new menu with no default option. In subsequent years, though prices changed significantly, consumers had their previously chosen plan as a default option so did not have to make an active decision. Second, after the menu change the firm offered three plan options with the exact same medical benefits (including their network of medical providers and covered services) but different financial characteristics. These three options range from a comprehensive plan with high premiums and low consumer cost sharing to a consumer driven health plan that has low premiums but high consumer cost sharing. The menu change and default structure give us the ability to observe choices in periods with and without switching costs while the relative plan features and detailed medical information allow us to accurately assess what plan values should be independent of switching costs.

Preliminary analysis of the data reveals evidence of large switching costs and a high degree of adverse selection. We present two simple tests to illustrate high switching costs. The first studies the behavior of new entrants to the firm over several choice periods. New entrants always have zero switching costs in plan choice since they are enrolling in a plan for the first time. We reveal that subsequent cohorts of new entrants, though almost identical demographically, make quite different choices during years when one cohort is new with zero switching costs and the previous cohort has a default option and positive switching costs. Specifically, the new cohort with zero switching costs makes choices that reflect current prices, while the older cohort's choices reflect previous prices and choices. Our second test takes advantage of the fact that, in some periods, it is possible for certain consumers to enroll in a plan that is completely dominated by another option.[3] We show that when a plan becomes dominated over time as a result of price changes, the majority of consumers who previously chose that plan, when it had better value, continue to enroll in that plan when it is their default option, after it becomes dominated. In addition to testing for switching costs, we use claims data from just prior to the menu change to show evidence of significant adverse selection after the menu change. Conditional on being enrolled in the same plan prior to the menu change, consumers who chose more comprehensive coverage after the change had approximately double the expenditures prior to the menu change relative to those who chose one of the less comprehensive options after the change.

While these tests reveal that switching costs and adverse selection are important, to precisely measure these effects and understand the impact of counterfactual policies we develop a structural choice model that jointly quantifies switching costs, risk preferences, and health risk. We are un-

---

[3]This means that for *any* realization of total health expenditures, one plan will deliver better financial value than another.

aware of any previous work that is able to estimate these important consumer choice elements together, likely due to high data requirements. In the model, consumers make choices that maximize their expected utility over all plan options conditional on their risk tastes and health risk distributions. When consumers are allowed to default into their previously chosen plan, switching costs reduce the utility of alternative options relative to the status quo option. We allow for heterogeneity in risk preferences so that we have the richest possible understanding of how consumers select plans.[4] Though not our primary focus, our estimates of risk aversion are interesting in their own right as we are aware of only one previous paper (Cohen and Einav (2007)) that quantifies a distribution of risk preferences in a non-experimental setting.

To precisely model consumer decision making, it is essential to accurately measure the health risk consumers perceive at the time of choice. We develop a novel methodology to predict individual specific ex ante distributions of medical expenditures that builds on the work of Carlin and Town (2007) and Abaluck and Gruber (2009). This is something that the vast majority of previous studies are unable to do because of a lack of detailed information on medical utilization. In our data, for each individual and each dependent, we observe every medical claim incurred at the same level of detail that the insurer has. We develop a framework that uses this data in conjunction with sophisticated predictive software developed at Johns Hopkins Medical School to predict each individual's distribution of medical expenditures for the upcoming year at the time of plan choice. The software takes the previous year's medical diagnoses and claims for each employee as an input and outputs a measure of predicted future resource utilization based on the types of medical claims incurred by each individual.[5] We use this predictive output as an input into a novel cost estimation framework that generates family-plan specific ex ante distributions of out-of-pocket expenditures. Assuming rational expectations, these distributions then characterize the uncertainty faced by each family for each plan. We note that for the cost model, we assume that there is no moral hazard (zero price elasticity of medical expenditure) and no cost relevant private information, an assumption that we provide substantial evidence in support of taking advantage of variation in coverage levels caused by the menu change.[6]

Our choice model estimates reveal that switching costs are high: an employee covering no dependents will on average forgo savings of up to $1,570$ dollars to remain enrolled in a default option from the previous year (the analogous figure for an employee covering dependents is $2,507$). This is a substantial amount as the average employee premium paid in the data is approximately $2,100$, while average total employee spending on medical care plus premiums is approximately

---

[4]There is a large body of previous work suggesting that there may be *advantageous* selection by consumers into health plans if very risk averse individuals who choose comprehensive insurance have low health risk (see e.g. Cutler, Finkelstein, and McGarry (2008)).

[5]For example, an otherwise healthy 35 year old male who breaks his arm and spends $10,000 will have much lower predicted resource utilization in the upcoming year than a 35 year old male with diabetes and hypertension who spent $10,000 in the past year.

[6]Our evidence on this point suggests that, even if moral hazard and private selection are present on some dimensions, their impact will be minimal and this assumption will not materially affect the results of the choice model. Further, since our initial health status measures are based on a year prior to the menu change, when all consumers were enrolled in a similar plan, these measures are not biased by potential moral hazard and private selection effects.

$4,500. Our estimated risk preferences reveal that the population is, on average, moderately risk averse without substantial heterogeneity on this dimension.

We apply these estimates to study a counterfactual policy experiment where consumers are provided with information that reduces their switching costs. We allow the impact of the policy to range from no reduction in switching costs to a complete elimination of switching costs. First, we conduct a partial equilibrium analysis to illustrate what the impact of the information provision policy would be if we viewed it from the perspective of the prior literature where there is no impact on market prices or adverse selection. We find that when the policy completely eliminates switching costs (implying rational and frictionless choice in all periods) improved consumer choices lead to a 10% increase in consumer welfare. The measure of consumer welfare we focus on is the mean change in the certainty equivalent value for each family as a result of the policy change divided by the mean employee expenditure on premiums. The 10% figure is for the population as a whole: for people who actually switch plans as a result of the policy the welfare gain is 19% while in this setup it is necessarily 0% for those who do not switch.

When we conduct the full equilibrium analysis, where individual choice improvements from the policy change can impact market prices, the results are quite different. In this analysis we assume a pricing structure that is identical to the way employee premiums are determined in the firm we study. The total premium (firm paid plus employee paid) equals the average cost of the enrollees in a plan during the previous period, conditional on the number of dependents, plus a small administrative fee. The firm subsidizes employees based on this total premium with an amount equal to an income dependent percentage of the lowest cost plan, implying that consumers pay the full marginal cost of more comprehensive insurance. Prices in the initial period when the firm changes its menu of plans are set equal to those the firm actually chose, which we observe. We simulate multiple years of pricing beyond the end of our dataset (holding health and demographics constant) to assess the long run implications of the policy change. This pricing environment is very similar to that in Cutler and Reber (1998) and Einav, Finkelstein, and Cullen (2009), and bears some similarity to the health insurance exchanges proposed in recent health care reform proposals.

We find that the improvement in consumer choices resulting from the policy significantly increases adverse selection in the full equilibrium setting. In the benchmark simulation where switching costs are as estimated in the data, enrollment in the most comprehensive insurance plan decreases by 25% over a four year period beginning at the time of the menu change. Once the policy intervention occurs and switching costs are zero, the market for comprehensive insurance almost disappears completely as enrollment decreases by 81% over the four year period. This marked decrease in enrollment occurs in conjunction with large relative price increases for comprehensive coverage resulting from that fact that only the sickest employees remain due to adverse selection. In the benchmark case, prices for comprehensive coverage increase by approximately 25% over the four year span while, with the policy intervention, they increase by over 100% for some dependent coverage tiers (and by a significant amount for all tiers). Strikingly, the policy intervention now *reduces* consumer welfare by 6% reversing the positive result from the partial equilibrium case. The

welfare of those who switch as a result of the inervention still increases, by 13%, while those who do not switch experience a 26.7% decline in welfare. In addition, the welfare of sicker consumers decreases slightly more than that of healthier consumers, while families are more adversely affected than individuals.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 provides some background on the institutional setting we study and describes the data. Section 4 presents preliminary evidence that switching costs and adverse selection are important in our environment. Section 5 describes the choice model and cost framework. Section 6 presents the model results. Section 7 analyzes the impact of a counterfactual information provision policy experiment that reduces switching costs. Section 8 concludes and discusses future research avenues.

2   Literature Review

2.1   Adverse Selection & Choice in Insurance Markets

There is a substantial recent literature that quantifies the welfare impact of adverse selection in health insurance markets. This work builds on the seminal theoretical advances in Akerlof (1970) and Rothschild and Stiglitz (1976). In the process of doing so, this literature has made significant advances in the structural modeling of health plan choice and utilization. Carlin and Town (2007) assess the welfare impact of adverse selection with a detailed panel data set on plan choice and utilization from a large employer. They are the only other paper we are aware of that uses the Johns Hopkins predictive medical software package to estimate individual specific medically motivated health risk measures. They conclude that adverse selection based on observable unpriced health information is large but that it leads to a very small welfare loss because consumer demand for insurance is price inelastic. Their model predicts very low insurance price elasticities because consumers have large and persistent unobserved plan preference heterogeneity. We believe that switching costs are a likely alternative explanation for why they observe substantial choice persistence and that choice persistence due to switching costs has very different policy implications than choice persistence due to unobserved plan preferences. When persistence is due to unobserved plan preferences, policy interventions designed to reduce switching costs will not change consumer choices or impact market equilibrium. Also, our model predicts that consumers are more price elastic in 'active' choice periods and less price elastic in 'passive' choice periods while their model predicts consumers are always price inelastic. They are unable to identify switching costs in their data because they do not observe a choice period where consumers must make new, 'active' choices as we do.

Cardon and Hendel (2001) study whether adverse selection based on private health information exists in a health insurance market. They perform a methodologically sophisticated empirical analysis that models choice and utilization behavior in a static setting with exogenous variation in the premiums consumers face. Bundorf, Levin, and Mahoney (2008) study the impact of different premium pricing rules on adverse selection and market efficiency. They find that government and

employee pricing policies lead to selection that reduces allocative efficiency by between 2-11%. Lustig (2008) finds substantial welfare losses associated with adverse selection in the market for privatized Medicare plans. None of these studies possess detailed health information at the consumer level, a panel of consumer behavior, or variation in the plan menu faced by a given population, making it impossible for them to study consumer choice and utilization at the level of detail that we do. Finally, Einav, Finkelstein, and Cullen (2009) develop a simple econometric specification to study the welfare consequences of adverse selection using detailed health plan choice and utilization data provided by a large employer. They develop a novel theoretical model describing the role of adverse selection and moral hazard in insurance markets and their estimates of this model imply that the welfare loss from adverse selection is small and costly to rectify.

More generally, Einav, Finkelstein, and Levin (2009) survey the empirical literature on insurance markets that structurally models choice foundations and firm behavior. They describe two distinctive types of choice models (i) those that build directly on expected utility theory and (ii) those that are closer to traditional discrete choice analysis and specify utility for a given option as a function of consumer and contract characteristics. The latter approach requires weaker assumptions on consumer utility but limits the ability to uncover parameters of interest such as risk preferences. Our approach falls within the former category as we directly model expected utility where uncertainty derives from the health risk distributions we estimate. This section of the literature is thin since the data requirements are higher and the methods less developed relative to the characteristic based approach. The leading example of other work in this area is Cohen and Einav (2007) who estimate risk preferences from data on car insurance deductible choices. They find a low level of median risk aversion, similar to our results, but document higher levels of heterogeneity in risk aversion than we do. Einav, Finkelstein, and Schrimpf (2009) use a similar setup to conduct welfare analysis of UK annuity markets. Cardon and Hendel (2001) is the one paper on health insurance mentioned above that falls into this line of research while Carlin and Town (2007), Bundorf, Levin, and Mahoney (2008), and Lustig (2008) take the more traditional characteristic based approach. Sydnor (2008) is another study that explicitly models choice under uncertainty while Bajari, Hong, and Khwaja (2006) and Finkelstein and Poterba (2008) are other examples of research in this area that use the characteristic based approach. Chiappori and Salanie (2000) survey the early work on testing for asymmetric information in insurance markets that this literature builds on.

Einav, Finkelstein, and Levin (2009) also highlight that this literature ignores choice frictions such as switching costs, which can have a significant impact on policy analysis. They note that such micro-foundations may be especially important in insurance markets where products are complex to evaluate. Our paper is the first work we are aware of that structurally estimates switching costs (or any departure from the standard choice paradigm) in a model with choice under uncertainty or a setting where adverse selection is a possible concern.

There is a significant literature studying adverse selection and plan choice in health insurance that does not structurally model choice or health risk foundations. Cutler and Reber (1998) study the trade-off between more adverse selection and greater competition that occurs when consumers

pay a higher marginal price to obtain more comprehensive insurance. They analyze a change in the subsidy policy at Harvard in the 1990's where the university went from heavily subsidizing comprehensive insurance to making consumers pay the entire marginal cost of that plan relative to cheaper plans. They find that this price change causes an adverse selection 'death spiral' where the most comprehensive plan disappears from the market entirely. They then show that the negative welfare impact of adverse selection is more than offset by the positive benefit of increased competition among insurers resulting from more price sensitive consumers. They do not possess medical data at the individual level nor do they observe a time period when all consumers make an 'active' plan choice. As a result, they are unable to estimate choice foundations such as switching costs and risk aversion which are important for any counterfactual policy analysis. Their analysis shows that an adverse selection 'death spiral' can occur while our estimates make it possible to study when one *will* occur for a variety of different policies affecting the level of switching costs, consumer subsidy level, or menu of available plans.

Cutler, Lincoln, and Zeckhauser (2009) study factors influencing the movement of individuals across health plans. They use detailed plan enrollment and medical data from employees of the state of Massachusetts to study the interaction between (i) adverse selection (ii) adverse retention (the tendency of the sick to remain in the same plan) and (iii) aging in place (plans with an older population increase in costs). They reveal that these factors influencing plan transitions significantly impact the evolution of market equilibrium. This work is notable in that it illustrates the way consumer choice behavior interacts with plan pricing in a dynamic environment. The firm pricing assumptions the authors make (plans price according to lagged average costs) are similar to those we make to determine pricing under counterfactual policies that reduce switching costs. Our research takes advantage of the plan menu change and forced re-enrollment to delve deeper into consumer choice and health risk foundations. Our more precise characterization of switching costs, risk preferences and health risks then allows us to analyze specific counterfactual policy experiments in a manner that their paper does not. In other related work Cutler, Finkelstein, and McGarry (2008) present empirical evidence that the degree of adverse selection in an insurance market depends on the correlation between risk preferences and health in addition to just the distribution of health. This reveals the importance of estimating risk preferences in conjunction with health risk in our setting. Strombom, Buchmueller, and Feldstein (2002) study health plan switching behavior as a function of age, health status, and job tenure and find that variation in these demographics impacts the price elasticity of demand. They are unable to precisely quantify these effects in a choice model, in part because they lack detailed medical data on enrollees. Tchernis, Normand, Pakes, Gaccione, and Newhouse (2005) study the relationship between dynamic changes in health status and plan switching and reveal that people who are more healthy are more likely to switch out of comprehensive insurance, exacerbating adverse selection. Fang, Keane, and Silverman (2008) provide some evidence for advantageous selection in privatized Medicare and document the cognitive ability of seniors making plan choices as one source of this selection.

Finally, within the context of our medical cost model we provide evidence consistent with a

low level of price elasticity of demand for medical services (moral hazard). Kowalski (2009) studies moral hazard using an instrumental variables strategy that relies on the fact that an injury to one family member shifts the marginal price of medical care for other family members. She finds a price elasticity of -2 for higher quantiles of the distribution of medical expenditures. The RAND health experiment conducted in the 1970s is still considered to be the standard in the literature. Consumers were randomized into health plans with different marginal prices and consumed whatever medical care they desired from that point forward. Estimates based on this study conclude that the price of elasticity for medical services is approximately -.2 (see Manning, Newhouse, Duan, Keeler, and Leibowitz (1987) or Newhouse (1993)).

2.2 Switching Costs & Information Provision

There is a growing literature studying the role of switching costs in consumer choice. Madrian and Shea (2001) investigate inertia in the choice of 401(k) savings plans. The authors use a switch in the default policy at a large firm to show that (i) 401(k) participation is significantly higher under automatic enrollment and (ii) many employees retain the default contribution rate and fund choice even though few employees hired before automatic enrollment picked those outcomes. In their environment, 401(k) participation is a one time decision and choice inertia is measured with respect to the default option. In our setting, employees first make an 'active' choice where no default option is allowed, and then are allowed to default into the plan they previously chose in subsequent periods. Our results add to those found in 401(k) choice by revealing that switching costs are high in an environment where consumers are allowed to default into a option they actively chose in a prior period. Esteves-Sorenson (2009) documents inertia in the choice of Italian television programs while Jones (2009) shows inertia in the choice of payroll deductions for income taxes.

There are a few studies that model switching costs within a structural choice framework. Dube, Hitsch, Rossi, and Vitorino (2008) study switching costs in the context of loyalty for grocery products. They include switching costs as an additive linear term in the utility for the previously chosen product and find that loyalty has a significant impact on choice. They show that high product loyalty has important implications for optimal brand pricing. While the way they model switching costs is similar to us, their context is quite different as there is no actual default option, uncertainty, or asymmetric information. Furthermore, they have no initial period without loyalty to identify preferences independent of switching costs. Shcherbakov (2007) estimates a dynamic model of consumer choice with switching costs in the television cable market. He argues that in an environment with switching costs, consumer should be forward looking and take into account future switching costs and market features in current decisions. We agree with this assessment in general but estimate our current specification with myopic consumers with an eye toward extending the model to a dynamic environment. One of the difficulties in our setting is that, in the absence of known future prices changes, we must consider several different assumptions on subjective expectations for future prices. Goettler and Clay (2009) estimate switching costs in a dynamic model of choice between grocery delivery plans. These costs allow the model to be consistent with the fact that

no consumers in the data switch plans within the firm despite the fact that delivery costs would be lower after doing so. Farrell and Klemperer (2007) and Klemperer (1995) survey a substantial theoretical literature on the impact of switching costs on market equilibrium. They reveal that switching costs can lead to a substantial 'lock in' effect through which firms can gain significant market power by accumulating market share. While our work focuses on the choice impact of switching costs and studies counterfactual pricing outcomes in a simplified firm setting, it will be interesting to consider the impact of the high switching costs we find within a richer dynamic firm model. To our knowledge, none of the previous work on switching costs considers the more complex impact they have in a market with adverse selection.

There are two studies that investigate choice inadequacy in the privatized Medicare market for prescription drug plans (Medicare Part D). Abaluck and Gruber (2009) use detailed linked plan choice and prescription drug utilization data to show that elderly consumers have a difficult time correctly assessing the value of 40+ drug plan options. They perform a counterfactual where they reduce the number of choice options and suggest that doing this in a targeted way can improve welfare. Importantly, they analyze this intervention within a partial equilibrium context so cannot comment on how this policy to improve consumer decisions impacts risk selection and plan premiums. Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2008) perform an experiment analyzing the impact of targeted information provision to seniors choosing Medicare Part D plans. The authors find that giving consumers personalized information, linked to their health status, on which plan is best for them improves their choices and welfare in a partial equilibrium setting. Again, this study does not investigate the consequences that this intervention has on adverse selection and plan pricing. However, this work does suggest that information provision can effectively improve consumer decisions, something that we will take as given in our counterfactual.

Sunstein and Thaler (2008) provide further evidence that 'choice architecture' can significantly impact the quality of consumer decisions. The authors illustrate that altering the framework within which consumers make decisions can change (sometimes drastically) choice outcomes.[7] Furthermore, they show that the effects of such interventions are more pronounced in settings where product choice is complicated, as in health insurance markets. This work provides a foundation for our counterfactual policy experiment that assumes information provision can reduce the level of switching costs we observe. The authors do not consider the impact of 'nudges' on equilibrium in a market with adverse selection but do analyze 'nudges' in such markets in a partial equilibrium context where we show the welfare implications may be quite different.

3   Data

We study data on health plan choice in an employer setting typical of the American health care system. In 2009, 59.3% of all individuals in the United States (159 million people) received insurance through their employer or the employer of a family member. Employer provided insurance is

---

[7]Potential interventions to change consumers' choice environments are, for example, (i) supplying consumers with targeted information and (ii) changing the type of default structure.

attractive relative to insurance purchased on the open market because employer health benefits are tax-deductible, implying that premiums can be paid for in pre-tax dollars. Still, the cost of employer provided health insurance is substantial and rising: in 2009 the average annual premium was $4,824 for single coverage and $13,735 for family coverage (a 131% increase over the analogous average 1999 premium). Employers typically subsidize the premiums paid by employees. In 2009, single employees contributed 17% of the their overall premium costs while families paid 27% of their total.[8] There is thus a significant amount of money at stake for consumers in this market, implying that the issues of choice inadequacy and adverse selection could lead to significant welfare losses.

3.1   Data Overview

We observe all the plan choices and health claims made by every individual at a large employer with approximately 9,000 employees. This includes detailed medical and cost information for each claim for each individual. We performed the data acquisition during an extensive process onsite at the firm. We worked with the firm's HR department to obtain plan choice and employee demographics linked at the individual level. In addition, they provided us with detailed information on the health plan characteristics and insurance provision structure at the firm over the time period we study. Finally, we worked with the firm's insurance administrator to acquire a database of all medical claims incurred by every individual enrolled within a subset of the plans offered by the firm. Specifically, the firm was self-insured for health plans covering approximately 60% of its population and owned the claims data for these plans. In the description that follows, the PPO options are the self-insured plans.[9] We merged this detailed utilization data with the choice and demographic information at the individual level to construct our final dataset.

The level of detail we possess on plan choice, demographics, and medical utilization for each individual is rare. For each employee we observe age, sex, gender, zip code, tenure with the firm, number and relationship of dependents, and month of entry/exit from the firm. We observe income for each employee grouped into one of five tiers where a higher numbered tier represents more income.[10] We can identify each dependent separately within the demographic and medical data. Lastly, we create a variable that indicates whether or not an employee is involved in a quantitatively sophisticated job or is a high level manager within the firm.[11] In addition to observing plan choices, plan premiums and detailed plan characteristics we observe several other contemporaneous decisions. For each employee we observe the dental and vision plan choices in each period, whether or not they enroll in a flexible spending account (FSA) or health savings account (HSA), and the contributions they make to these accounts. For each employee and each covered dependent enrolled in a PPO

---

[8]The figures in this paragraph are taken from Kaiser Family Foundation (2009).

[9]Large firms self-insure in order to avoid costly state insurance mandates. When a firm self-insurers, it takes on the health risk usually passed to the insurer but outsources plan administration and network construction to a traditional insurance company.

[10]The tiers change slightly from year to year but span approximately $45,000 each, up to the last tier which covers all earners above a given threshold. We do not reveal the exact tier levels to protect the privacy of the firm. The firm uses these five tiers for premium pricing purposes as described soon.

[11]We observe some information on job status that we use to construct these variables. Further detail is available upon request.

option we have detailed medical utilization data. For each individual we observe the payment and medical information for each claim. On the payment side we observe details like deductible paid, coinsurance paid, copayment, insurer paid, total billed charges, and whether the claim is in or out of network. On the medical side we observe ICD-9 diagnostic codes (which divide the universe of diagnoses into 25,000 categories), CPT procedure codes, and the medical provider. In addition, we observe medically relevant aggregations of diagnoses, procedures, and provider specialties derived and used by the insurer.

The data contain 14,248 employees making yearly plan choices for at least some of the period 2004-2009. Many employees cover dependents, bringing the total covered lives we observe to 25,214. There is entry and exit from the sample each year such that there are about 9,000 employees (17,000 covered lives) in any given year. Each year approximately 1,500 employees waive insurance coverage (likely because they or their spouse have an alternative source of coverage). The roughly 7,500 employees (covering 14,000 lives) who remain choose one plan from a menu of five health plans that changes over the sample period. Unless they enter the firm mid-year, employees choose a plan with a set premium level in November for the following calender year. Over that year individuals demand medical services; the interaction of services undergone and plan-characteristics such as deductible and coinsurance determine how much an individual pays for medical expenses out-of-pocket each year, in addition to their plan premium. The first column of Table 1 describes the demographic characteristics of the entire sample.

[Table 1 about here.]

From 2004 to a year denoted t_{-1} the employer offered four HMO plan options (more vertically integrated with providers) and one PPO option (less integration).[12] Each of the five plans had a different network of providers, different contracts with providers, and different premiums and cost-sharing formulas for enrollees. For the year after t_{-1}, denoted t_0, the firm changed the menu of plans. From year t_0 onward the firm continued to offer five plans but reduced the number of HMOs to two and introduced three new PPO options. This plan structure remained intact through 2009. After the menu change, the HMOs still had different providers and cost sharing rules relative to the set of PPOs. However, the three new PPO plans introduced in the menu change had the same provider network, covered medical services, and contractual treatment of providers. The only difference between these plan options were the premiums and plan financial characteristics that determine the mapping from total claim expenditures to medical out-of-pocket expenditures (e.g. deductible, coinsurance, and out-of-pocket maximums).

Further, the three PPO options after the menu change also had the same provider network and covered services as the PPO option before the menu change and, again, only differed from that plan based on financial characteristics. We denote the PPO option before the menu change as PPO_{-1}. We denote each PPO option after the menu change by its individual deductible so that it is clear throughout which plans are more comprehensive (pay a higher percentage of medical claims). The

---

[12]The firm requested we not reveal the year of the menu change in order to protect its identity.

'best' or most comprehensive plan is PPO_{250}, the middle plan in terms of cost sharing is PPO_{500}, and the plan with the highest cost sharing is PPO_{1200}. PPO_{1200} is paired with a health savings account option that allows consumers to deposit tax-free dollars to be used later to pay medical expenditures. This may lead to some degree of differentiation for this plan relative to the other two.[13] We denote the HMO options available throughout the entire period as HMO_1 and HMO_2 and those available only before the menu change as HMO_3 and HMO_4.

In addition to changing the menu of plans for year t_0, the firm forced consumers to leave their old plans and make an 'active' plan choice from the new menu for this year. The firm told employees that there would be no default for year t_0, implying that they would either have to actively select a plan or lose out on this valuable benefit. During the enrollment period, employees were continuously contacted and reminded of this.[14] Further, unless an employee was previously enrolled in HMO_1 or HMO_2 they had to choose a new plan by definition implying they could not just choose their previous option without actively considering the new menu. In years prior to and after t_0 employees were allowed to default into their previously chosen plan option without taking any action. This was true even though in many cases plan prices changed significantly from one year to the next. This variation in the structure of the default option over time is a unique feature of the data set that makes it especially well suited to study switching costs because, for each employee, we observe at least one choice where switching costs can be present and one choice where they are not.

Table 2 lists the financial characteristics of the PPO options. The two higher deductible PPO options have double the coinsurance rate, which is the marginal price of medical care post-deductible, as PPO_{250}. Out-of-pocket maximums indicate the maximum amount of medical expenditures that an enrollee can pay post-premium in a given plan.[15] In our sample, out-of-pocket maximums are larger when the plan deductible is larger and vary according to an enrollee's income level within each plan. PPO_{250} and PPO_{500} have identical copayment (flat fee) structures for pharmaceuticals and physician office visits while these services apply to the deductible and coinsurance in PPO_{1200}, implying very different high-end marginal prices for these services in that plan. All plans charge 50% coinsurance on mental health claims once the plan deductible is met. Finally, each plan has distinct out of network characteristics, which we do not list here since only 2% of overall claims are out of network.

[Table 2 about here.]

Figure 1 illustrates total employee plan payments (including premiums and out-of-pocket costs) as a function of total in-network hospital and physician claims for PPO_{250} and PPO_{500} in year t_0

---

[13]This kind of plan is known as a 'high-deductible health plan' or 'consumer driven health plan'. The health savings account is considered to be a hassle to use for medical payments relative to standard payment mechanisms but may also be seen as a beneficial extra retirement account since funds can be withdrawn tax-free for any purpose after the age of 65.

[14]Eventually, though they were not told this ahead of time, the 50 employees that did not actively elect a given plan were enrolled in PPO_{500}.

[15]Technically this applies only to certain expenses, something that we take into account in our model. For example, mental health expenses do no apply to the out-of-pocket maximum in any plan.

(the chart applies to low income families). These two plans are identical except for this mapping, so this figure encapsulates the total financial benefit of choosing one plan over the other. Throughout our analysis we assume that (i) premiums are in pre-tax dollars and (ii) medical expenses are in post-tax dollars.[16] When a family has zero total claims, the vertical intercept of Figure 1 is the premium contribution for each plan. For positive claims, a family pays all expenses until it reaches its deductible after which it pays the coinsurance percentage of marginal claims. Eventually, if a family spends enough, it hits its out-of-pocket maximum, after which it pay zero percent of all remaining medical expenses.[17] The chart reveals that in year t_0, healthy families should choose PPO_{500} while sicker families should choose PPO_{250}.

[Figure 1 about here.]

3.2 Plan Pricing

Each plan offered by the firm has a total premium and employee premium contribution in each year. The total premium is the full cost of insurance while the employee premium contribution is the premium the employee actually pays after receiving a subsidy from the firm. Who determines the total premium depends on the plan. For the HMO options the insurer determines the total premium. For the self-insured PPO options the firm determines the total premium in conjunction with advice from the plan administrator, who is a large insurer. In our data, the total premiums for all plans reflect the average cost of their enrollees in the previous period plus an administrative loading fee that is between 15-25%. For the HMOs, competition from other insurers forces total premiums to be close to average cost. For the PPOs the firm could theoretically choose any total premiums it wants to reflect its distributional aims. In our sample firm total premiums are not arbitrary and follow a lagged average cost rule.[18] For each PPO plan in each year, the total premium set by the firm equals the average plan cost of its enrollees in the previous period, conditional on the dependent coverage tier.[19] The firm claims it uses this policy in order to (i) evaluate and run each plan as an independent unit and (ii) make employees bear the cost of comprehensive insurance on the margin to ensure efficient plan selection.

---

[16]In reality, medical expenses may also be in pre-tax dollars since individuals can pay medical expenses with FSA and HSA contributions which are pre-tax. However, since approximately 25% of the population enrolles in these accounts, we make this assumption. Ultimately, we could make the tax treatment of medical expenses individual specific since we observe savings account contributions. In order to convert premiums into pre-tax dollars we multiply premiums by an income-contingent combination of state and federal marginal tax rates using the NBER TAXSIM data.

[17]In the plans we study, each family member technically has his or her own deductible and out-of-pocket maximum. On top of these, the family has an aggregate deductible and aggregate out-of-pocket maximum that limits what can be paid toward the individual deductibles and out-of-pocket maximums respectively. This chart assumes that expenses are allocated proportionally across family members. The individual and family limits are taken into account in estimation.

[18]This is true except for year t_0 when the new plans have no previous data to assess average cost. The employer set a total premium in this period based on some other metric, possibly what their expectation of average costs acress the plans would be. We delve into the firm's pricing process in more detail in our counterfactual section.

[19]For year t_0 and after the firm had four dependent coverage tiers: single, spouse only, child(ren) only, and spouse plus child(ren). Before t_0, the firm had two dependents tiers: single and family.

Given total premiums for each plan, the firm sets a target of subsidizing 76% of total employee premium payments. Since the HMO options are not self-insured, by law the firm must provide the same subsidy to all employees conditional on the number of dependents covered (the premiums can and do vary based on this). For the PPO options, the firm can subsidize plans in any way it wants to as long as it doesn't condition the subsidy on factors linked to individual health risk. The firm conditions PPO subsidies each year on an employee's income tier, in addition to the number of dependents, because of equity considerations. The firm provides PPO subsidies as a fixed percentage of the total premium for PPO_{1200} where the fixed percentage depends on income tier and the total premium is contingent on number of dependents. This fixed subsidy is decreasing with income. As a result of this policy, employees deciding between the PPO options after the menu change pay the full marginal cost of PPO_{250} and PPO_{500} premiums relative to PPO_{1200}. We note here that since total premiums are group-rated (by law) and determined by lagged average cost, making employees pay the marginal premium for comprehensive insurance makes significant adverse selection likely. In this scenario, it is likely that sicker individuals will be willing to pay the marginal cost of more insurance, healthier individuals will not, average costs for comprehensive insurance will increase, and relative premiums will increase.

Table 3 illustrates employee premium contributions in year t_0 and year t_1 for the single and family (spouse plus children) coverage tiers. There is a noticeable decrease in premiums for PPO_{500} relative to the other plans from t_0 to t_1 coupled with a slight increase in the premium for PPO_{250}. This premium movement is due to adjustment based on t_0 average costs for each of these plans. Employee premium payments are much higher in t_1 for PPO_{250} relative to PPO_{500} even though all other plan characteristics remain exactly the same as in t_0. As a result, the choice setting in year t_1, when most employees had a default option and switching costs, is quite different than that in t_0, when the forced re-enrollment occurred.

[Table 3 about here.]

3.3 Plan Choices

Table 4 presents the plan choices of all employees for years t_{-1}, t_0, and t_1. In year t_{-1}, before the menu change, approxiamtely 50% of employees that enroll in a plan enroll in PPO_{-1} while the other half of the population enrolls in one of the four HMO options. After the menu change, in years t_0 and t_1, about 60% of employees enrolled in a plan enroll in one of the three available PPO options. For those employees who enroll in a PPO after the menu change, approximately 55% choose PPO_{250} at t_0 while 45% do so at t_1. In each year, approximately 15% of employees waive coverage.

[Table 4 about here.]

Table 5 reveals that there is substantial horizontal differentiation between the PPO and HMO nests of plans. Specifically, we show that if an employee ever chooses a PPO option, they are

unlikely to switch to an HMO option at any point in time. This is important to our upcoming analysis, which focuses on the sample of employees who (i) were enrolled in PPO_{-1} at t_{-1} and (ii) continued to be enrolled in a PPO option at the firm for at least the next two years through t_1. The top panel in table 5 studies the choice behvaior of all employees at the firm who were enrolled in any plan in both of the years t_{-1} and t_0. It is clear that, when the menu of plans changed for t_0, most employees in PPO_{-1} moved to one of the new PPO options while most employees enrolled in an HMO at t_{-1} still re-enroll in an HMO at t_0. The bottom panel presents the analogous chart for all employees at the firm enrolled in a plan both in years t_0 and t_1. The vast majority of t_0 PPO enrollees who switch plans at year t_1 choose another PPO option at t_1. Since employees who ever choose a PPO are very likely to *always* choose a PPO, even if they switch plans, the unmodeled availability of HMO options will not bias the results of analysis restricting the set of options to just the PPO plans.

[Table 5 about here.]

4 Preliminary Results

4.1 Switching Costs

The raw data provide substantial evidence of switching costs in plan choice when consumers are given a default option. This section presents two separate pieces of evidence supporting this.

First, we study the behavior of new employees at the firm. New employees are unique because in any year they join they must have zero switching costs. This is because (i) they have no health plan default option and (ii) they were not previously enrolled in any health plan within the firm. Thus, employees who are new in year t_0 have zero switching costs in that period and positive switching costs when choosing a plan for year t_1 (along with the rest of existing employees) while employees who are new in year t_1 have zero switching costs at t_1. If switching costs exist, and the profile of new employees is similar from year to year, we would expect the year t_1 choices of employees who were new in year t_0 to reflect both persistence toward their t_0 choices and year t_1 prices while we would expect t_1 choices by t_1 new employees to reflect only year t_1 prices.

Table 6 reveals those who were new enrollees in year t_0 do in fact make substantially different choices in t_1 than new employees at t_1. The table reveals that these cohorts of new entrants have almost identical demographic characteristics, which suggests that in the absence of switching costs these two populations should make similar choices in year t_1 when they are both at the firm. The data in the table show that, in fact, these two groups make significantly different decisions in year t_1. Moreover, these choices differ in exactly the way a model with switching costs would predict: the choices of year t_0 new employees exhibit significant persistence in t_1 relative to year t_0 choices while the choices of year t_1 new employees strongly reflect the price shift in favor of PPO_{500} in year t_1. For example, 21% of year t_0 new employees choose PPO_{250} at t_0 and 23% choose PPO_{500}. At t_1, 20% of this population chooses PPO_{250} and 26% choose PPO_{500} only a small change in total share for each plan. However, in year t_1 only 11% of year t_1 new employees choose PPO_{250} while

43% choose PPO_{500}. This indicates that t_0 employees made active choices at t_0 and only adjusted slightly to large price changes at t_1, due to high switching costs, while t_1 new employees with no t_1 switching costs made active choices at t_1 reflecting the current prices.

[Table 6 about here.]

Our second piece of evidence takes advantage of a unique situation caused by the combination of plan characteristics and plan pricing changes. Figure 2 combines the chart comparing plan out-of-pocket expenses for year t_0, Figure 1, with the analogous chart at t_1. The t_1 chart looks quite different due to the significant price change in favor of PPO_{500}. In fact, at t_1, for the low income family category depicted, a family would be better off choosing PPO_{500} for *any* level of medical expenditures. PPO_{250} is similarly dominated in t_1 for four of O

individual's health risk, risk preferences, and level of switching costs. Even though it is clear from these examples that switching costs are present, quantifying them along with health risk and preferences is essential to understanding how the entire population will choose plans, and how the market will evolve, in counterfactual settings. This motivates the structural choice and utilization model we present in the next section.

4.2 Adverse Selection

Before we present the main econometric framework, we provide evidence that adverse selection is present in the data we observe. Here, we study selection based on the combination of observable but unpriced health information as well as private health information. The next subsection studies adverse selection based only on private information. Here, we analyze the subset of the population that (i) was enrolled in PPO_{-1} at t_{-1} and (ii) continued to be enrolled in a PPO plan at the firm for at least the next two years through t_1. Restricting the sample in this way allows us to compare selection across the PPO options with the same population over all three years. Column 3 in table 1 describes this sample, which is slightly older and slightly richer than the full sample. As noted earlier, there is almost no crossover from the nest of PPOs to the nest of HMOs during this time period so analysis within the set of PPO options can stand on its own.

Table 8 presents evidence of adverse selection against the most comprehensive PPO option, PPO_{250} once the menu change occurs at t_0. The table shows the level of t_{-1} claims for individuals enrolled in each of the PPO options from t_{-1} to t_1. We use t_{-1} claims in every year since all families in the sample were enrolled in the same plan (PPO_{-1}) that year, implying that total claims are an 'apples to apples' measure of health expense risk that is not confounded by moral hazard. Risk-based selection in the sample is striking, as employees who chose PPO_{250} in both year t_0 and year t_1 had almost double the median and mean of t_{-1} total medical claims relative to enrollees in the other two PPO options.

[Table 8 about here.]

Despite the large price change from year t_0 to year t_1 the pattern of selection barely changes from t_0 to t_1. The initially high level of selection reveals that consumers initially chose plans based on health risk while the lack of movement in selection over time implies that individuals did not update their selection based on health risk in years where they could default into their previous option. Thus, high switching costs likely reduce adverse selection relative to what it would have been if everyone made active plan choices in each period. This insight motivates our counterfactual studying the impact of an information provision policy that reduces switching costs.

4.3 Moral Hazard & Selection on Private Information

The change in the menu of available plan options for year t_0 provides a natural experiment that makes it possible to analyze whether or not consumers who selected less comprehensive coverage at t_0 had systematically lower total claims in that year than at t_{-1}. Evidence that consumer expenditures

are systematically lower based on subsequent coverage choice could be indicative of moral hazard (positive price elasticity) or risk selection based on health information not represented in prior expenditures.[21] We do not view this analysis as rigorous evidence for or against moral hazard and private selection, but do believe that it provides support for our cost model assumption of no moral hazard or selection based on unobservable information.[22]

We investigate the same subset of employees as in the adverse selection analysis. This includes those employees who were enrolled in PPO_{-1} at t_{-1} and enrolled in any PPO option after the menu change through at least year t_1. It is important for our analysis that only one PPO option was available prior to t_0; this ensures that we observe one year when all consumers have the same price incentives. We split the population into two groups (i) employees who enrolled in PPO_{250} in year t_0 and (ii) those that enrolled in one of the two less comprehensive PPO options in t_0. We refer to the first group as the 'control' group and the second as the 'treatment' group. The marginal price of medical care in PPO_{250} is similar to that of PPO_{-1} before the menu change implying that employees who choose PPO_{250} face the same medical price incentives in both years. As a result, assuming a constant distribution of health in this group over the two years, the difference between expenses in this population at t_{-1} and t_0 should be entirely due to price inflation. Alternatively, the deductibles in PPO_{500} and PPO_{1200} are more than double those of PPO_{250} while, post-deductible, the marginal price of care (coinsurance) is twice as high in these less comprehensive plans. Further, families paid much higher rates on the margin for pharmacy and physician office expenses in PPO_{1200} than in all other plans. Thus, in the treatment group, some of the difference in expenses from t_{-1} to t_0 may be due to moral hazard or selection based on health information not linked to t_{-1} expenses. In the treatment population, we group PPO_{500} and PPO_{1200} together in order to increase sample size. When we analyze these two plans separately the results are unchanged.

Table 9 shows how the distribution of medical expenses changed over time for the control and treatment groups. Denote the claims of the control group in year t as κ_t^C and those of the treatment group κ_t^T. If the combination of moral hazard and selection on private information is relevant, we would expect:

$$\frac{\kappa_{t_0}^C}{\kappa_{t_{-1}}^C} > \frac{\kappa_{t_0}^T}{\kappa_{t_{-1}}^T}$$

Thus, we would expect inflation in medical expenditures to be the same across both groups but the treatment group to have a reduction in expenses relative to the control group due to moral hazard or selection based on private information.[23] In our sample, the results point in the opposite direction. For each portion of the distribution of expenditures, expenses increase by 10% more in

---

[21]Our analysis does not distinguish between moral hazard and selection based on private information at t_0. Simple theoretical models predict that the effects of both these phenomena will lead to a positive correlation between total claims and more insurance.

[22]In future versions of this work, we plan on doing robustness checks that estimate the cost model with minimal but positive bounds on the level of moral hazard.

[23]This is essentially a log difference in differences setup where one difference is with repect to claims over time and the second difference is for the control group relative to the treatment group.

the treatment group relative to the control group (20% vs. 10%).[24] When we consider the absolute changes in medical expenditures within each group, instead of the ratios, the same conclusion holds. Since the ratio moves in the opposite direction from what theory would predict, if the test is not misspecified then, for any statistical test, we could not reject the null hypothesis of no moral hazard or selection not based on prior observable expenses. We note again that this analysis is done to support our assumption of no moral hazard and private information selection in the cost model. Further analysis in future work is required to rigorously isolate these effects.

[Table 9 about here.]

We delve into the issue in more depth in Table 10, which analyzes the same ratio of year t_0 expenditures to year t_1 expenditures for the treatment and control group with finer groupings of medical diagnoses. We investigate the problem at this level in case there are specific categories of elective medical expenditures where moral hazard and time-varying selection are important. We study 20 aggregated medical diagnostic groupings derived by the insurance administrator from the finer ICD-9 codes. We study the ratio of median of t_0 total claims in each category relative to the median of t_{-1} claims, conditional on expenditures being greater than zero.[25] Of the 20 diagnostic groupings, the median of treatment group expenditures decreased by a higher percentage than the median of control group expenditures in only three aggregated diagnostic categories. The difference between these ratios is close to zero for seven groups and expenditures increase relatively in the treatment group for the remaining ten categories. This evidence suggests that the prevalence of moral hazard and selection based on private information is low across different classes of medical procedures as well as for aggregate claims.

[Table 10 about here.]

Finally, we formalize this analysis at the level of aggregate diagnosis groups. To do this, we construct ratios of expenditures over time within each diagnostic category for individuals within the treatment and control groups and run the following regression:

$$log(\kappa_{id,t_0}) = \delta_d + \alpha \ log(\kappa_{id,t_{-1}}) + \beta \ log(\kappa_{id,t_{-1}}) * \mathbf{1}_T + \epsilon_{id}$$

Here, $\kappa_{id,t}$ refers to the total claims of individual i in aggregated diagnostic category d in year t. δ_d is a category fixed effect and $\mathbf{1}_T$ is an indicator variable that equals one if a family is in the treatment group. We estimate this equation with a median quantile regression and find that β is negative, statistically significant at a 5% confidence level, and equal to 3.7% of α in magnitude. This implies that conditional on t_{-1} expenses, individuals in the treatment group reduced expenditures by 3.7% in t_0 relative to those in the control group. The full results from the mean and median

---

[24]We note that this chart looks at the ratio of distributional derivatives and not the distribution of ratios in the population. We feel that this is appropriate since individual-specfic ratios will be noisy on a year to year basis but the population distribution of health should be nearly constant.

[25]There is almost no movement in the number of claims in each category and each group.

regressions are presented in Table 11. This difference in differences result reveals that when we investigate expenditure ratios over time at the individual level there is evidence of a combined moral hazard / private information effect but one that has a very small magnitude. Since the estimated effect is small in magnitude, we believe that our cost model assumption of no moral hazard or selection based on private information does not bias the results of that model in any meaningful way.

[Table 11 about here.]

5 Empirical Framework

The preliminary results reveal that switching costs and adverse selection are both substantial. However, in order to understand the impact of counterfactual policies that reduce the level of switching costs, it is necessary to precisely measure all determinants of plan choice. To this end, we construct and estimate a structural model of consumer choice that quantifies switching costs, risk preferences, and health risk. Once we know these fundamentals, we are able to examine the equilibrium impact of counterfactual policies that reduce switching costs.

We restrict our estimation sample to employees (and covered dependents) who are (i) enrolled in PPO_{-1} in t_1 and (ii) continue to enroll in a PPO option at the firm through at least t_1. This sample is described in the third column of table 1 and is the same one used in our preliminary adverse selection analysis. This sample construction is advantageous for three reasons. First, these restrictions imply that for years t_0 through t_2 we observe a previous year of detailed claims data for each individual in the sample. This allows us to use the same rigorous cost framework for the entire sample.[26] Second, the restriction to the set of PPO options implies that all plans in the choice set have the exact same provider network and covered medical services. As a result, our choice model does not need to take into account unobserved preference heterogeneity for medical plan benefits, which must be done in most prior plan choice studies. Finally, since each family must be enrolled in a plan from t_{-1} to t_1, we observe each family making a plan choice in a year when they have a default plan option and a year when they do not, a feature that aids our identification of switching costs. We note that, as shown previously, there is almost no movement from PPO plans to HMO plans over time, even at t_0, implying preferences for these nests are highly differentiated. This makes us confident that the set of PPO options at t_0 and after is an appropriate description of the choice set for this sample.[27]

Our empirical framework has two primary components (i) a choice model and (ii) a cost model. The cost model predicts family-plan specific ex ante distributions of medical expenditures for the

---

[26] If we included new entrants after t_{-1} or individuals enrolled in HMOs in the sample, we could use a less precise cost framework for these employees in the absence of detailed medical information. This could be based on future claims or demographics such as age and sex, similiarly to what is done in the rest of the literature when detailed claims information is not available.

[27] For example, table 5 reveals that, of the 2,757 employees enrolled in PPO_{-1} in year t_{-1} who also enroll in a plan at t_0, only 85, or 3%, choose an HMO option at t_0.

upcoming year, which are used as inputs into the choice model to quantify expense risk. We first present the choice model, taking health expense risk as given, and then describe how the cost model predicts distributions of expenditures.

5.1 Choice Model

Our choice framework quantifies risk preferences and switching costs conditional on the family-plan-time specific distributions of health expenditures output by the cost model. We denote these expense distributions by $F_{kjt}(\cdot)$, where k denotes a family, j denotes one of the three PPO options available after the menu change, and t denotes one of three time years after the menu change (t_0 through t_2). We describe the precise way that $F_{kjt}(\cdot)$ is generated later in this section.

We assume that families have rational expectations with respect to their health status and that these expectations confirm to the cost model output $F_{kjt}(\cdot)$. Each family has latent utility U_{kjt} for each plan in period t. In each time period, each family chooses the plan j that maximizes U_{kjt}. We use what Einav, Finkelstein, and Levin (2009) call a model of 'realized' utility and assume that U_{kjt} is the v-NM expected utility of each family:

$$U_{kjt} = \int_0^\infty f_{kjt}(OOP)u(OOP, P_{kjt}, \mathbf{1}_{kj,t-1}, W_k, X_k, Y_k, H_k)dOOP$$

Here, $u(\cdot)$ is the v-NM utility index and OOP is a realization of medical expenses from $F_{kjt}(\cdot)$. Employees pay premium contributions P_{kjt} that, as described earlier, depend both on covered dependents and on income. $\mathbf{1}_{kj,t-1}$ is an indicator of whether the family was enrolled in plan j in the previous time period. W_k and X_k denote family-specific wealth and income respectively. [28] [29] Finally, Y_k is a category indicator for whether or not the employee covers any dependents while H_k is an indicator of whether the family is high cost. We determine whether or not a family is high cost by a threshold that depends on family size.[30] This latter variable is included to proxy for the fact that most families with very high expenses are likely to choose PPO_{250} whether it is the best plan for them or not. It is possible that these families assume that, because they have high expenses, they should always choose the most comprehensive insurance option (a heuristic which is supported in the raw data).

We assume that individuals have constant absolute risk aversion (CARA) preferences implying that for a given consumption level x depending on $(OOP, P_{kjt}, \mathbf{1}_{kj,t-1}, W_k, X_k, Y_k, H_k)$:

$$u(x) = -\frac{1}{\gamma}e^{-\gamma x}$$

---

[28]We don't observe family wealth separately from income. In our main specification, preferences over the menu of options does not depend on wealth so we are agnostic about this variable.

[29]We assume that income is constant over time since we observe only small movements over the sample for a given employee and link income to time-invariant risk preferences in the model.

[30]For example, this indicator has value one if a single individual spends more than \$15,000 or a family spends more than \$30,000 in any year in the sample. Both Y_k and H_k are taken to be constant over time as they index time invariant intercepts in the model.

Here, γ is a risk preference parameter. As γ increases, the curvature of u increases and the decision maker is more risk averse. The CARA specification implies that the level of absolute risk aversion $\frac{-u''(\cdot)}{u'(\cdot)}$, which equals γ, is constant with respect to the level of x.[31] We also studied a constant relative risk aversion (CRRA) specification that yielded similar results to our primary framework.

In our model the level of consumption for each draw of OOP from $F_{kjt}(\cdot)$ depends on multiple factors. A family's utility $u(\cdot)$ for any given OOP draw for plan j at time t is:

$$u(\cdot) = \frac{-1}{\gamma_k(X_k)} e^{-\gamma_k(X_k)(W_k - P_{kjt} - OOP + \eta(Y_k)\mathbf{1}_{kj,t-1} + \delta_k(Y_k)\mathbf{1}_{1200} + \alpha_j(Y_k)H_k + \epsilon_{kjt})}$$

Here, γ_k is a family-specific risk preference parameter. Since we don't observe γ_k, it is estimated as a random coefficient where the distribution of γ depends on income X_k. Thus, conditional on income, we allow for unobserved heterogeneity in risk preferences. η is a switching cost that is constant conditional on Y_k. δ_k is an unobserved family-specific plan intercept for PPO_{1200} that is estimated as a random coefficient ($\mathbf{1}_{1200}$ is an indicator for $j = PPO_{1200}$). On average we expect δ_k to differ from zero because the health savings account option offered exclusively through PPO_{1200} could cause significant hassle costs or provide an extra benefit in the form of an additional retirement account.[32] α_j measures the intrinsic preference for plan j for a high-cost family as a function of family status. ϵ_{kjt} represents a family-plan-time specific idiosyncratic preference.

This specification assumes a particular functional form that has implications for how switching costs are interpreted. Here, at time t we assume that the utility for the plan chosen at $t-1$ increases by η when that plan is the default option at t (this is true for all contingencies in our sample). This is similar to the approach taken in the literature (see e.g. Dube, Hitsch, Rossi, and Vitorino (2008) or Shcherbakov (2007)). This framework implies that, on average, for a family to switch at time t they must prefer an option other than their default option by \$$\eta$ more than their default. This assumes that switching behavior is linked to the relative utility between plans at t. In our final section, we discuss an alternative specification where switching behavior can also be independent of relative plan utilities. We view this as a promising avenue for future work that can help distinguish whether switching costs are related to relative differences in plan utilities or due to a fixed cost of re-optimization (or both).

In addition, this model of switching costs assumes consumers are myopic as they do not make dynamic decisions. In theory, switching costs are a dynamic phenomenon since choice in the current period impacts the structure of utility for plans in future periods. We believe that since price changes are not signaled in advance (and likely not expected to change in any specific direction) the model with myopic consumers is a plausible description of our environment. Our final section discusses future work where forward-looking consumers make dynamic choices.

Finally, our specification includes preference factors that are constant over realizations from the

---

[31]This implies that wealth does not impact relative plan utilities. As a result, it drops out in estimation.

[32]Apart from these rational considerations, this intercept could incorporate a shift in the distribution of the error term ϵ_{kjt} for PPO_{1200}. For example, if employees used a choice heuristic that steers them away from the least comprehensive plan, this will be reflected in a negative value of δ_k.

health risk distributions, such as switching costs and plan intercepts, in the v-NM utility index instead of adding these terms linearly outside of the expected utility formulation. We take this approach so that these preference parameters are estimated in constant dollar units: if taken outside the expectation over health risk a constant η would mean different things for people with different γ_k. In the CARA benchmark model presented above, these constant terms can be treated as a constant shift to wealth across plans and do not have different implications over different OOP realizations.

5.2 Cost Model

We described the choice model in the previous section taking $F_{kjt}(\cdot)$ as given. This section describes how we generate the family-plan-time specific ex ante distributions of health expenses that are used as an input into the choice model. Appendix A provides a more formal treatment of this material.

Given the evidence presented earlier on moral hazard and selection based on private information, we maintain the assumption that there is no moral hazard or selection based on private information in this section. Our preliminary results imply that even if these factors are relevant they will be of a small magnitude that will not strongly impact the cost and choice model estimates. [33] We note that, because our cost model combines very detailed medical utilization information with sophisticated medical diagnostic software, most information that is traditionally unobservable in health studies will be observable in our model. This makes additional selection based on private information much more unlikely than it would be in a model that uses coarse demographics or aggregate health information to measure health risk. [34] [35]

In order to incorporate as much detail as possible, our model predicts health risk at the individual level and then aggregates these predictions to the family level. In order to model out-of-pocket expenditures across the menu of plans for a given individual, it is necessary to predict multiple categories of medical claims. This is necessary because the mapping from total claims to out of pocket medical expenditures depends on how total claims are divided up over different types of claims. For example, in all plans, hospital and physician (non office visit) expenses apply to deductible and coinsurance while in some of the plans pharmacy expenditures are subject to fixed copayments per drug. We capture this heterogeneity by dividing total claims into four mutually exclusive categories of claims and estimating the joint distribution of these claims for each individual. The four distinct claims categories we consider are **(i)** hospital and physician **(ii)** pharmacy **(iii)** mental

---

[33] In an independent study Cardon and Hendel (2001) find no evidence of selection based on private information. In future work, as a robustness check we could study the alternative assumption that moral hazard is bounded below some quantity, and analyze the impact this has on our results relative to the model presented here.

[34] Pregnancies, genetic pre-dispositions, and non-coded disease severity are possible examples of private information that could still exist. Carlin and Town (2007), whose cost model is done with detailed medical information, also argue that significant residual selection is unlikely. In future work, we plan to include a pregnancy dummy at the time of choice in order to control for selection on this dimension. This is possible since we can 'back date' pregnancies to the time of plan choice in the claims data.

[35] It is also possible the individuals in our sample know *less* about their future risk profile since we use medical diagnostic software to map past conditions into future health risk.

health and **(iv)** physician office visit.[36] Each of these claims subdivisions contributes uniquely to the plan-specific mappings from total claims to out-of-pocket expenditures. Table 12 describes the mean and median of total claims (plan paid plus family paid) for each of these categories across the different plan options. We are unaware of any previous work that reconstructs the mapping from individual-level total claims to plan-specific out-of-pocket expenditures at this level of detail.

[Table 12 about here.]

In the majority of models studying consumer choice and utilization in health insurance, health risk is either modeled based on **(i)** demographic variables such as age and gender and/or **(ii)** aggregated medical cost data from past or futures years. These approaches are useful approximations when detailed medical data are not available, but imprecisely characterize a given individual's information set at the time of plan choice. For our detailed choice analysis, it is essential to have a more accurate picture of health risk at the time of plan choice. We model utilization at a higher level of precision by using our detailed claims data in conjunction with sophisticated medical predictive software developed at Johns Hopkins University.

The software we use is called Johns Hopkins ACG, version 8.2, where ACG stands for adjusted clinical grouping. This program is one of the most widely used and respected risk adjustment and predictive modeling packages used in health care.[37] For our purposes, the software predicts the level of future medical expenditures for each individual based on past medical claims information. Importantly, this prediction heavily involves the actual past medical diagnoses for each individual. For example, a 35 year old male who spent $10,000 on diabetes last year would have higher predicted future health expenses than a 35 year old male who spent $10,000 to fix a broken arm. As a result, our predictions of future expenditures for each individual are very precise because we get close to their exact information set at the time of plan choice. In addition to treating medical conditions separately, the software analyzes the medical relevance of combinations of medical conditions to predict future expenses.

Since the description of the cost model setup and estimation is lengthy, we briefly summarize the model here and present the details of this framework in an appendix. First, we use the Johns Hopkins ACG software to map each individual's previous claims and demographic data into (i) a measure of total predicted future resource utilization and (ii) a measure of predicted future pharmacy utilization. Second, based on the output of the Johns Hopkins model we group individuals with similar predicted risk into cells for each of the four claims categories. Within each category, we estimate a distribution of expenditures for the upcoming year for each cell based on the ex post claims realizations of claims for that group in that category. The decision to use a cell-based approach, rather then a more continuous model, has costs and benefits. One cost is that all individuals within a given cell for a given type of claims are treated identically. A benefit is that our

---

[36]In our cost model appendix, we discuss how we arrived at this division of claims. We omit out-of-network payments and emergency room copayments from this mapping because they do not materially affect its accuracy. Note that we still include these claims in the cost model in the hospital and physician category.

[37]For more information on this program visit $http://www.acg.jhsph.edu/html/AboutACGs\_whois.htm$.

method produces local cost estimates for each individual that are not impacted by the combination of functional form and the health risk of medically different individuals. Since our sample size is large enough to create fairly fine cell groupings, we believe that the benefit of localized flexible estimates across the four claims categories outweighs the cost of treating individuals in the same cell identically.

Once we have an estimated future claims distribution for each cell and category, we generate individual-specific joint distributions of claims across all four categories. It is necessary to construct joint distributions because ex post health shocks are empirically correlated across the four categories. We use the empirically observed rank correlations across claims categories for each total health status cell to combine each individual's marginal claims distribution in each of four cells into one joint distribution. We then construct a detailed mapping from a realized vector of claims across the four categories to out-of-pocket expenditures for each plan j. This mapping takes into account the exact manner in which each total claims in each category map to out-of-pocket expenditures, as described in table 2.[38] The more detailed analysis in the appendix uses the actual claims data to verify the accuracy of this mapping. We map each individual's joint distribution of total claims across the four categories into a plan-specific distribution of out-of-pocket expenditures, and then combine individual distributions at the family level taking into account family-level plan characteristics (e.g. family deductible limit). The final cost model output is the set of family-plan-time-specific distributions of out-of-pocket expenditures $F_{kjt}(OOP)$ for all k, j, and t. This is used as an input into the choice model to quantify family-specific expectations of future medical expenditures at the time of plan choice.

5.3 Choice Model Identification

Our choice model identification strategy exploits the features of the data highlighted in section three. Two unique features of the data are essential to identify switching costs. These are (i) the forced re-enrollment of consumers into new health plans created by the menu change at t_0 and (ii) the significant price changes that occur from t_0 to t_1 (and t_1 to t_2 as well). At t_{-1} all employees in our sample were enrolled in PPO_{-1}, an option that was not available after the menu change at t_0. Thus at t_0 there is no default option or existing option that consumers were enrolled in previously. So, at t_0 all plans are on even ground and there are no switching costs. For year t_1 choice switching costs are relevant because individuals have a default option (their previous plan). Since prices at t_1 are significantly different from prices at t_0, if switching costs are low we should observe many consumers in the sample switching plans. For example, if switching costs are zero, choices at t_1 should reflect the stable preferences estimated at t_0 applied to the environment with new prices (and potentially different health expense distribution from the cost model). If switching costs are high, then choice will change little even though the stable preferences estimated at t_0 predict they otherwise would. Thus, conditional on the estimates of the stable preference parameters (e.g. risk

---

[38]For example, this mapping takes into account plan deductible, coinsurance, specific copayments, and out-of-pocket maximums along with other plan characteristics.

aversion, plan intercepts) identified at t_0 and the health expense distributions estimated in the cost model, the variation between choices at t_0 and t_1 identifies switching costs. Though new entrants in any year provide another source of identification for switching costs (since they have zero switching costs when they arrive) we do not include these consumers in the choice model because we require prior claims data to estimate ex ante cost distributions.[39]

The distribution of the family-specific risk aversion parameter, γ_k, is identified separately from switching costs using the choices employees make at year t_0, when there are no switching costs. We take advantage of a novel feature of the dataset, namely, that we observe individuals choosing over three health plans with the exact same medical benefits and provider network when there must be zero switching costs. Conditional on our cost model output, each family's choice at t_0 is a choice between different risky gambles (with the additional random intercept for PPO_{1200} representing differentiation based, potentially, on the health savings account feature). Thus, unlike with most datasets, our estimates of risk aversion are not confounded with unobserved heterogeneity in preference for different plans due to differences in medical networks or covered services. The risk aversion coefficients are identified separately from the plan intercept for PPO_{1200} because choices between PPO_{250} and $PP0_{500}$ at t_0 are not affected by switching costs or this intercept. Thus, bounds on the risk aversion coefficients implied by choices between these options combined with our estimation assumptions identify the distribution of risk parameters.

5.4 Choice Model Estimation

To estimate the choice model we make assumptions on the distributions of unobserved parameters γ_k and δ_k. We study two specifications for risk preferences, one where γ_k is normally distributed and one where it is log normally distributed. In both versions the mean of this risk heterogeneity shifts based on income X_k. The normal specification is:

$$
\begin{aligned}
\gamma_k(X_k) &\rightarrow N(\mu_\gamma(X_k), \sigma_\gamma^2) \\
\mu_\gamma(X_k) &= \mu + \beta(X_k)
\end{aligned}
$$

The random intercept for PPO_{1200}, δ_k, is distributed normally with mean $\mu_\delta(Y_k)$ and variance $\sigma_\delta^2(Y_k)$. We assume that both the mean and variance of this intercept depend on whether the employee covers dependents to allow for maximum flexibility along this dimension.

Switching costs are constant conditional on Y_k. For simplicity in what follows we denote the switching costs for single employees by η_s and those for families by η_f. We use the same notation for the plan specific high-cost intercepts α_j, such that single employee values are α_{js} and family values are α_{jf}. We normalize the value of $\alpha_{PPO_{250}}$ to zero since this is a relative measure. Finally, we assume that the error terms ϵ_{kjt} are mean zero and drawn from i.i.d normal distributions with

---

[39]In future work, we may be able to incorporate new entrants in the estimation sample using more coarse measures of ex ante health related to consumer demographics and ex post claims.

unknown variances σ_{ϵ_j}. We normalize the value of $\epsilon_{kPPO_{250}t}$ to zero since this is a relative metric.

The model is estimated via simulated maximum likelihood since this has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. The likelihood is computed for a *sequence* of choices for each individual since switching costs imply that the likelihood of a choice made in the current period depends on the choice made in the previous period. The maximum likelihood estimator selects the parameter values that maximize the similarity between actual choices and choices simulated with the parameters.

We simulate Q draws from the distribution of health expenditures, F_{kjt}, for each family, plan, and time period. These draws are used to compute plan expected utility conditional on all other preference parameters. We simulate Z draws from the distributions of risk preference and plan preference random coefficients, conditional on the set of parameters:

$$\theta \equiv (\mu, \beta, \sigma_\gamma, \mu_\delta(Y_k), \sigma_\delta(Y_k), \alpha_j(Y_k), \sigma_{\epsilon_j}, \eta(Y_k)).$$

We denote θ_z one draw from these parameters and θ_Z the set of all Z simulated draws. For each θ_z we use all Q health draws to compute family-plan-time-specific expected utilities. We simulate choices in each period using a smoothed Accept-Reject Simulator with the form:

$$Pr(j = j^*) = \frac{(\frac{\frac{1}{-U_{kj^*t}}(\cdot)}{\Sigma_J \frac{1}{-U_{kjt}}(\cdot)})^\tau}{\Sigma_{\hat{j}}(\frac{\frac{1}{-U_{k\hat{j}t}}(\cdot)}{\Sigma_J \frac{1}{-U_{kjt}}(\cdot)})^\tau}$$

Here, U_{kjt} is the discrete utility for each plan and each family in each time period, conditional on the simulated preference draw z and Q health draws. Theoretically, conditional on these draws we would want to just pick the j that maximizes U_{kjt} for each family. However, doing this leads to a likelihood function with flat regions because for small changes in the estimated parameters the discrete choice made does not change. This is why smoothing is necessary. In our specification, since all utilities are negative with CARA, the highest negative utility, or the utility with the lowest absolute value is the best choice. This is why in our smoothing function $\frac{1}{-U_{kj^*t}}$ is the object of interest since as this object becomes larger an option is more preferred. If this is higher for j^* than for another j, an individual is more likely to choose j^*. By choosing τ to be large, an individual will always choose j^* when $\frac{1}{-U_{kj^*t}} > \frac{1}{-U_{kjt}} \forall j \neq j^*$. In estimation, we make τ large so this smoothed Accept-Reject simulator becomes almost identical to the true Accept-Reject simulator that chooses the actual utility-maximizing option in each case with probability one.

Denote any sequence of three choices as j^3 and the set of such sequences as J^3. In the limit as τ grows large the probability of a given j^3 will either approach 1 or 0 for a given simulated draw z and family k. This is because for a given draw the sequence (j_1, j_2, j_3) will either be the sequential utility maximizing sequence or not. This implicitly includes the appropriate level of switching costs by conditioning on previous choices within the sequential utility calculation. For example, under θ_z a choice in period two will be made by a family k only if it is optimal conditional on θ_z, other

preference factors, and the switching costs implies by the period one choice. For all Z simulation draws we compute the optimal sequence of choices for k with the smoothed Accept-Reject simulator, denoted j^3_{zk}. For any set of parameter values θ the probability that the model predicts j^3 will be chosen by k is:

$$\hat{P^{j^3}_k}(\theta, F_{kjt}, X_k, Y_k) = \Sigma_{z \in Z} \mathbf{1}[j^3 = j^3_{zk}]$$

Let $\hat{P^{j^3}_k}(\theta)$ be shorthand notation for $\hat{P^{j^3}_k}(\theta, F_{kjt}, X_k, Y_k)$. Conditional on these probabilities for each k, the simulated log-likelihood value for parameters θ is:

$$SLL(\theta) = \Sigma_{k \in K} \Sigma_{j^3 \in J^3} d_{kj^3} ln \hat{P^{j^3}_k}$$

Here d_{kj^3} is an indicator function equal to one if the actual sequence of decisions made by family k was j^3. Then the maximum simulated likelihood estimator (MSLE) is the value of θ in the parameter space Θ that maximizes $SLL(\theta)$. These are the estimated parameters we present in the next section.

6   Choice Model Results

Table 13 presents the results of the choice model. We present results from the benchmark model where the risk parameter γ is assumed to be normally distributed and from an alternative specification where γ is log normally distributed.

[Table 13 about here.]

The estimated switching costs are large in magnitude for both single individuals and employees covering dependents. On average, when a single employee can default into a previous option, he forgoes up to $\$1,571$ in expected savings from an alternative option to remain in that default plan. The analogous amount for employees who cover dependents is $\$2,507$. These figures are estimated at a high level of precision and are stable across the two specifications. There are several potential underlying sources for these switching costs including (i) time and hassle costs (ii) re-optimization costs and (iii) inattention resulting from a status quo bias.[40] We address some potential ways to distinguish between these sources in the extensions section. The high level of switching suggests that policies that reduce switching costs, which we study in the next section, could have a significant impact on consumer choices and market outcomes. The switching costs standard errors (and those for all parameters) are small, indicating a high level of statistical precision (e.g. the 95% confidence interval for the individual switching cost parameter is [1304, 1834]).

Table 14 interprets the estimates of absolute risk aversion from both specifications. Our analysis of these coefficients follows that found in Cohen and Einav (2007). We interpret these coefficients by

---

[40]Another discussed source of switching costs, having better information about your current plan, is similar but not identical to re-optimization costs.

determining the value X that would make an individual indifferent between inaction and accepting a gamble with a 50% chance of gaining \$100 and a 50% chance of losing \$X. Thus, a risk neutral individual will have $X = 100$ while a very risk averse individual will have X close to zero. With normal heterogeneity, X is 93.6 for the mean / median individual implying a moderate amount of risk aversion. X is 88.9 for the 95th percentile of γ and 86.6 for the 99th so preferences are focused in a region of moderate risk aversion. The log normal specification yields similar results with the exception of a larger tail of individuals who have high risk aversion (X is 78.1 for the 95th percentile and 60.5 for the 99th). The results presented in Table 14 are for employees in the median income tier (tier three). Our estimates in Table 13 reveal that risk aversion is slightly increasing in the level of income though these results do not significantly alter the interpretation in Table 14.[41] Finally, Table 14 compares our results to those found in the literature. Our mean estimates lie in the middle of the range of those from previous studies while the log normal results are of the same order of magnitude as those in Cohen and Einav (2007), the only other non-experimental study we are aware of that estimates risk preference heterogeneity.

[Table 14 about here.]

The remaining estimates in Table 13 indicate that, conditional on monetary considerations, (i) there is a strong distaste for PPO_{1200} and (ii) very high cost individuals are more likely than the remainder of the population to choose PPO_{250}, conditional on the monetary benefit of doing so. For example, our estimates reveal that the average employee covering dependents prefers either PPO_{250} or PPO_{500} by \$5148 over PPO_{1200} independent of monetary considerations. The variance of \$2148 on this plan intercept implies that a few families will consider PPO_{1200} on close to neutral ground while most will have a strong distaste for this plan. There are at least two reasons why consumers might have a distaste for PPO_{1200}. First, the health savings account tied to this plan may lead to significant hassle costs. Though consumers choosing this plan can opt out of this account, if they do so they will then be forced to pay all out of pocket medical expenditures in post-tax dollars, leading to a significant financial loss. Second, consumers choosing a plan based on heuristics may be wary of choosing a more complicated plan that is the most risky option in the choice set. We suggest an alternative specification in our discussion of future work where employees may choose a plan based on heuristics instead of an explicit calculation of costs and benefits.

Finally, our estimates predict that single employees with high medical costs, on average, prefer PPO_{250} by \$758 over PPO_{500} and by \$2,212 over PPO_{1200} conditional on the PPO_{1200} intercept and all other costs and benefits (for families, these parameters are \$1,655 and \$3,506 respectively). The innate preference of high cost employees for the most comprehensive plan offered likely stems from the fact that these employees assume that, because they have high utilization, the top tier option *must* be the best for them. Then, without an explicit consideration of costs and benefits, these employees will be more likely to choose PPO_{250} all else equal.

---

[41]The positive relationship between income and risk aversion is possibly there because (i) higher income employees have a heuristic that causes them to select higher coverage and (ii) we don't estimate heterogeneity in plan intercepts with respect to income(doing so would be a way to test this hypothesis).

7   Counterfactual Policy Analysis: Information Provision

There is a general consensus in the policy debate on the design and regulation of insurance markets that helping consumers make the best plan choices possible in each period is unequivocally the right course of action, regardless of the specifics of the environment. Currently, most U.S. national health reform proposals include the creation of a health insurance exchange, which is a heavily regulated market in which private insurers compete for a specified risk pool of consumers.[42] Advocates of health insurance exchanges cite the need for consumers to make well-informed, optimal choices. In an essay proposing an economic foundation for exchanges, Enthoven, Garber, and Singer (2001) write:

> "Exchanges must make available comparative information on plan
> benefits, pricing, quality measurement, quality improvement
> initiatives, and other aspects of plan performance in an effort
> to help members make informed, high-value choices." (p. 158)

There are two primary rationales for promoting better choice in insurance markets (and markets in general). First, improving choices helps the match quality between consumers and health plans leading to increased consumer welfare conditional on prices and the market environment. Second, improving choices promotes competition among insurers and, in the long run, may help promote more efficient insurers. However, these positive effects are only part of the story. This analysis reveals that improving choices in markets with adverse selection may exacerbate selection leading to lower overall welfare and unforeseen distributional consequences.

In this study, consumers make poor plan choices over time because of switching costs. After initially making informed decisions, consumers are allowed to default into their previously chosen plan. Over time, prices change significantly but switching costs stop consumers from adjusting to these large market changes. In this section we study the impact of a policy intervention designed to improve consumer choices over time by reducing switching costs. In each period when consumers can default into their previous option, they are provided with targeted information that reduces their level of switching costs by some amount between zero and the full level of estimated switching costs. We quantify the consumer welfare impact of this policy in (i) a partial equilibrium setting where the price of insurance does not change as a consequence of selection and (ii) a full equilibrium setting where plan prices change to reflect the new risk profile of employees enrolled in the different options. We find that the policy intervention improves consumer welfare by 10% of the overall premiums paid in the partial equilibrium case but *reduces* this welfare by 6% of total premiums paid in the full equilibrium setting. This analysis assumes that the cost of the information provision policy is zero. In reality, a positive cost for this policy would reduce these welfare numbers.

---

[42]This risk pool is usually a combination of employees of small firms and unemployed individuals. Some proposals suggest allowing everyone to purchase through the exchange.

We study the impact of information provision by assuming that the improved consumer knowledge reduces switching costs by some value Z. Our analysis investigates the choice and welfare impact for Z in between 0, when there is no information provision, and $\hat{\eta}$ when all switching costs are removed. This approach assumes that there are a range of possible information intervention exercises, some of which are more effective than other in reducing switching costs. In the limit as Z increases to $\hat{\eta}$ one could imagine essentially selecting a plan for each consumer based on their stated preferences. For this analysis we restate the expected utility of family k for plan j at time t as an explicit function of premiums P_{kjt} and switching cost $\eta_k - Z$:

$$U_{kjt}(P_{kjt}, \eta_k - Z) = \int_0^\infty f_{kjt}(OOP)u(OOP, P_{kjt}, \eta_k - Z, \mathbf{1}_{j-1})dOOP$$

Here, for notational simplicity, the dependence of utility on factors besides switching costs and price is not explicitly stated since these are the only two choice relevant features that change as a result of the information provision exercise. The v-NM utility index with these prices and switching costs is:

$$u(OOP, P_{kjt}, \eta_k - Z, \mathbf{1}_{j-1}) = \frac{-1}{\gamma_k(X_k)}e^{-\gamma_k(X_k)(W_k - P_{kjt} - OOP + (\eta_k - Z)\mathbf{1}_{kj,t-1} + \delta_k \mathbf{1}_{1200} + \epsilon_{kjt})}$$

Thus, relative to the estimated choice model, consumers may make different choices because switching costs favoring the default plan option are reduced and because this reduction in switching costs may indirectly lead to a change in premiums. From this point forward we will use the expression $\hat{U}_{kjt}(P_{kjt}, \eta_k - Z)$ to denote the family-plan-time specific expected utilities under the parameters estimates from the model with normally distributed risk preferences provided in the previous section. A family chooses the plan j at time t that maximizes expected utility:

$$\max_{j \in J} \ \hat{U}_{kjt}(P_{kjt}, \eta_k - Z)$$

In both the partial and full equilibrium cases we analyze consumer welfare using a certainty equivalent approach (our approach is similar to that used in Einav, Finkelstein, and Cullen (2009)). For each plan j, the certainty equivalent is the monetary amount Q that makes a consumer indifferent between getting Q for certain and obtaining the risky payoff of enrolling in j. This welfare measure is useful because it translates the expected utilities, which are subject to cardinal transformations, into values that can be interpreted in monetary terms. Certainty equivalents are calculated based on utilities net of switching costs: if an individual chooses the plan he is defaulted into his utility for welfare purposes is $\hat{U}_{kjt}$ removing switching costs $\eta - Z$. Thus, welfare is computed as if employees always make active choices even though their choices are made taking switching costs into account.

Denote the certainty equivalent of plan j for a family k at time t as CEQ_{kjt}. In what follows, we'll describe a policy intervention by the amount Z by which it reduces switching costs. Conditional on Z, we denote these certainty equivalents CEQ_{kjt}^Z. By this metric, the absolute change in consumer

welfare for a given individual is the difference between the certainty equivalent of the choice j_z made after the policy intervention and the choice j made in the benchmark model:

$$\Delta CS_{kjt} = CEQ^Z_{kj_Zt} - CEQ_{kjt}$$

The mean change in consumer welfare in absolute terms at time t as a result of the policy intervention is:

$$CS_t = \frac{1}{\|K\|} \Sigma_k \ \Delta CS_{kjt}$$

In order to determine the percentage change in welfare at time t as a a result of policy intervention Z we divide the mean change in consumer welfare by three different measures derived from the benchmark case with switching costs. These three metrics are (i) the average premium paid in year t (ii) the average sum of premium and out of pocket medical expenditures at t and (iii) the average certainty equivalent of the plans consumers enroll in at t.

In the partial equilibrium case prices do not change from those observed in the data so changes to $\hat{U_{kjt}}(P_{kjt}, \eta_k - Z)$ and subsequent welfare metrics depend only on the impact that Z has on choices. In the full equilibrium setting, prices change as the reduction in switching costs causes consumers to switch across plans so our utility and consumer welfare measures depend on the change in prices as well as the new choices. We assume that premiums change according to the pricing rule used by the firm during the time period studied. The firm was self-insured for the PPO options so has full control over the total premiums assigned to each plan and the subsidies given to employees toward those premiums. Based on outside advice and internal policy decisions,the firm sets total premiums 'as if' it purchased plans from insurers who determined those premiums. As a result, in each year, the total premium for each plan option and each dependent coverage tier y, TP^y_{jt}, is the average plan cost for the previous year's enrollees in that group plus an administrative fee denoted L:

$$TP^y_{jt} = AC_{K^y_{j,t-1}} + L = \frac{1}{\left\| K^y_{j,t-1} \right\|} \Sigma_{k \in K^y_{j,t-1}} PP_{kjt-1} + L$$

Here, $K^y_{j,t-1}$ refers to the population of families in plan j at time $t-1$ and PP_{kjt-1} is the total *plan paid* amount by plan j for that family at $t-1$. Thus, total premiums for each plan in each year are determined by taking the plan's previous year average cost for each dependent coverage tier and adding an administrative loading fee.

The total premium TP^y_{jt} is the amount an employee in dependent category y enrolling in plan j would have to pay if they received no health insurance subsidy from the firm. In our setting, the firm subsidizes insurance for each employee as a percentage of the total PPO_{1200} premium based on income tier X_k. Denote this subsidy proportion as $S(X_k)$.[43] Then, for any dependent coverage tier y, the portion of the premium for plan j that the employee actually pays is:

---

[43]In our setting the subsidy rates for the five income tiers ordered from poorest to wealthiest are .97, .93, .83, .71, and .64.

$$P_{kjt} = TP_{jt}^y - S(X_k)TP_{PPO_{1200}t}^y$$

For PPO_{1200} the employee contribution is a fixed percentage of the total premium. For the other two options, employees pay the full marginal cost of the total premium relative to PPO_{1200}. This can alternatively be thought of as a lump sum subsidy tied to the total premium of PPO_{1200}. Making employees pay the entire marginal premiums for more comprehensive insurance is a standard regulatory feature in health insurance exchange proposals that is thought to lead to more efficient plan choices since consumers internalize price signals. However, high marginal prices also exacerbate adverse selection since plan price differences due to distinct health risk pools are magnified. This feature is important when analyzing the full equilibrium consequences of reducing switching costs since consumers are more price responsive once switching costs are lowered. Cutler and Reber (1998) study this tradeoff between improved price shopping and increased adverse selection in a similar environment where consumers pay the full marginal cost of more comprehensive insurance.[44]

We note that this pricing structure is nearly identical to that assumed in both Einav, Finkelstein, and Cullen (2009) and Cutler and Reber (1998) where the authors assume that health plans compete for employees in the firms they study. We believe that enriching the firm environment along several dimensions is an interesting avenue for future work, which we discuss further in our final section.

In addition, while our model describes how the firm will update prices in each period as a function of plan costs, it does not stipulate how the firm initially sets prices when there is no prior plan history (at t_0 in our sample). Clearly, in our case the firm sets t_0 prices below what the long run equilibrium prices are after consumers select into plans based on health risk. Since the market pricing rule depends on last period's average cost, the path of prices and the long-run stationary price in our environment depends on the prices set at t_0. In both the partial and full equilibrium setups we keep the same initial period prices set by the firm but note that this impacts the results; if the initial relative price of PPO_{250} was set too large compared to what the static equilibrium price would be then information provision could reduce adverse selection over time instead of exacerbate it. In this sense these counterfactuals illustrate the impact of information provision subject to these initial price conditions. We view determining why different initial pricing conditions may arise endogenously as an important topic for future work.

7.1 Partial Equilibrium Results

Our partial equilibrium simulation assumes that the prices we observe in the data are fixed but that our information provision exercise reduces switching costs by Z dollars. Consumers may select new plans, increasing or decreasing the level of selection based on risk, but this selection will not feed back into prices as it does in the full equilibrium case. We investigate behavior in the counterfactual setting for the years t_1 and t_2. Year t_0 choices and welfare are the same with or without the policy intervention because the forced re-enrollment implies that there are no switching costs at t_0. The

---

[44]In our environment where the firm is self-insured this subsidy policy is an effective way for the employer to push employees toward less comprehensive insurance.

benchmark of full switching costs is simulated from the parameter estimates as well in order to ensure consistency.

Table 15 presents the choice and average cost results from the simulation for $Z = 0$ (the benchmark case), $Z = \frac{\eta}{2}$, and $Z = \eta$ (complete removal of switching costs). The removal of switching costs helps consumer make better decisions as more consumers adjust to the price change in favor of PPO_{500} in both cases where switching costs are reduced. For example, at t_2 enrollment in PPO_{500} increases by 72% relative to the benchmark case when switching costs are reduced to zero (from 573 to 994). Almost all of these new enrollees caused by the policy intervention switch from PPO_{250}, which was initially more attractive but became much less so after the price change. Average costs for families (employee plus spouse plus child(ren)) in PPO_{250} increase as switching costs are reduced implying that the people who switch out of PPO_{250} as a result of the policy intervention are healthier than those that do not switch.[45] This suggests that in a full equilibrium setting, where plan prices adjust to reflect the new health risk profile, PPO_{250} premiums will become more expensive, potentially leading to even more selection against that plan in subsequent periods.

[Table 15 about here.]

Table 16 presents the welfare results of the policy intervention in a partial equilibrium context. The table analyzes the change in certainty equivalents resulting from the policy that reduces switching costs to 0 ($Z = \eta$). When we use total premiums as the baseline for welfare at stake, the policy intervention improves welfare by 8.6% in year t_1 and 10.3% in year t_2. This metric incorporates the percentage change in welfare for the entire population, including those individuals who switch plans as a result of the intervention and those that do not. In partial equilibrium, individuals who do not switch by necessity have zero welfare change. The welfare changes at t_1 and t_2 conditional on switching as a result of the intervention are 16.4% and 19.0% respectively. When the welfare baseline is the absolute value of certainty equivalent dollars at stake, the population welfare changes at t_1 and t_2 are 3.2% and 3.6%.[46]

[Table 16 about here.]

We note that in a partial equilibrium context, the policy intervention can only increase welfare since prices are unchanged and the policy simply helps consumers make better decisions when switching costs are high. This is the paradigm for all previous work on resolving choice inadequacy in insurance markets (see e.g. Kling, Mullainathan, Shafir, Vermeulen, and Wrobel (2008) and Abaluck and Gruber (2009)). In the behavioral industrial organization literature there is some work studying the equilibrium consequences of consumer choice inadequacy (see e.g. Gabaix and Laibson (2006) or Ellison (2006)) but this work does not study the context where choices directly impact costs as they do when adverse selection is a concern. We now turn to the full equilibrium simulations where prices change along with choices.

---

[45]The impact on average costs is similar across coverage tiers; families are selected as an example.

[46]Using the absolute value of the total certainty equivalent dollars at stake is necessary since expenditures are a negative consumption flow. Thus, the certainty equivalent is a negative number which is less than total expenditures.

7.2   Full Equilibrium Results

In the full equilibrium analysis prices change as consumers switch plans due to the policy interven-
tion. As a result, adverse selection may be exacerbated when switching costs are reduced because
healthy individuals who would have remained in comprehensive insurance as prices rise will now
switch, making the risk pool for that comprehensive plan sicker on average. In fact, the partial
equilibrium results suggest this will happen since the average costs of PPO_{250} enrollees rises as a
result of the information intervention. When prices can change, the welfare impact on employees
that do not switch plans is no longer zero: the change in their plan price resulting from risk selection
will increase or decrease their welfare. Similarly, for employees who do switch plans there will now
be two effects contributing to the welfare change (i) increased welfare from choosing the best option
conditional on prices (ii) a theoretically ambiguous welfare change from the price effect. For expo-
sitional purposes, the results we present in this section compare the case where $Z = 0$ to the case
where $Z = \eta$. As in the partial equilibrium case, interim values of Z lead to intermediate results
on other dimensions. Since we simulate prices as well as choices in this environment, we can study
time periods beyond the data set to investigate the long-run impact of the policy intervention. We
perform simulations for the years t_0 to t_6. As before, year t_0 choices are unaffected by the policy
intervention and the benchmark $Z = 0$ case is simulated to maintain consistency.

Figures 3 presents the time path of choices for PPO_{250} and PPO_{500} with and without the policy
intervention to reduce switching costs. In full equilibrium, the impact of the policy intervention
on the market share of PPO_{250} relative to PPO_{500} is noticeable. In the benchmark case where
there are significant switching costs over the six year period the market share of PPO_{250} declines
from 1237 to 793 while that of PPO_{500} increases from 458 to 882. After the policy intervention
reduces switching costs to zero, PPO_{250} enrollment declines all the way to 255 after six years
while PPO_{500} enrollment increases to 1416. Thus, the policy intervention to reduce switching costs
almost ends up eliminating comprehensive insurance (or causing a comprehensive insurance 'death
spiral' as discussed in Cutler and Reber (1998)). Relative to the partial equilibrium case at year
t_2, PPO_{250} has much lower enrollment after the policy intervention, further revealing that the full
equilibrium effect incorporating adverse selection leads to significantly different market outcomes
than the partial equilibrium analysis.

[Figure 3 about here.]

Figure 4 presents the time path of average costs and prices for single individuals enrolled in
PPO_{250} and PPO_{500} with and without the policy intervention. Not surprisingly, these two items
move closely in tandem since prices are linked directly to average costs. Prices are given for in-
dividuals in the middle income tier. The policy intervention has a noticeable impact on both the
prices and average costs of $PP0_{250}$ relative to PPO_{500}. When switching costs are reduced to zero,
the employee premium of PPO_{250} increases from $3,071$ for year t_1 to $5,390$ for t_6 while in the
benchmark model with switching costs the long run price is $3,929$. The long run price of $PP0_{500}$
without switching costs is $2,044$ while with them it is $1,733$. Thus, the employee premium for

PPO_{250} both increases substantially with the policy intervention *and* increases by a lot relative to the premium of PPO_{500}. This large long run price differential leads to the large reduction in market share for PPO_{250} depicted in Figure 3.

[Figure 4 about here.]

It is clear that the policy intervention to reduce consumer switching costs exacerbates adverse selection leading to very high prices and low market share for comprehensive insurance. Table 17 presents results on the welfare impact of the policy intervention. By all metrics, welfare decreases as incremental adverse selection occurs over the six periods studied. For year t_6, the policy intervention to reduce switching costs leads to a population welfare reduction of 5.9% using total premiums in that year as a baseline. Relative to previous work, this is a surprising result: helping consumers improve their own choices conditional on prices reduces overall welfare in the population. When we use total employee spending or total certainty equivalanet as the baseline, welfare decreases by -2.8% and -2.0% respectively. In addition, the table reveals that the policy intervention has distributional consequences. At t_6 employees who switch plans as a result of the policy have 13% higher welfare relative to the no intervention case. Conversely, employees who do not switch due in t_6 due to the policy change have a 26.7% welfare reduction. This is interesting to contrast with the partial equilibrium results where non-switchers have zero welfare loss by necessity. Our results also illustrate that the policy hurts both unhealthy and healthy individuals by a similar amount while causing a much larger welfare reduction for employees covering dependents than for single employees.

[Table 17 about here.]

8   Conclusions and Future Work

This paper used a unique data set to investigate the equilibrium consequences of a policy designed to improve consumer choices in a market with adverse selection. Simple tests based on the data alone reveal large switching costs and substantial adverse selection. We estimate a structural choice model that jointly estimates switching costs, risk preferences, and ex ante health risk in order to precisely quantify these important choice fundamentals. We use these estimates to study the welfare consequences of a counterfactual information provision policy that reduces switching costs and improves consumer choices. In a partial equilibrium context, where insurers do not adjust premiums as their risk pools change, the policy intervention increases consumer welfare by 10%. However, once we extend the analysis to a full equilibrium setting where premiums can change, we find that the policy intervention still improves choices but also exacerbates adverse selection. The combination of these effects leads to a 6% welfare loss. This result is surprising in light of previous work that only studies choice inadequacy in partial equilibrium, where policies to improve decision making can only increase welfare.

37

There are multiple possible avenues to improve the research presented here. On the demand side, there are at least three clear ways to enrich our choice framework. First, our model assumes that consumers are myopic and do not consider future switching costs when making current choices. It would be fruitful to model dynamic consumer behavior in order to (i) test if consumers behave dynamically and (ii) determine if allowing for dynamic choice changes our results. Second, our model assumes that consumers have rational expectations of their future health expenditures and fully understand the mapping between claims and out-of-pocket expenditures in each health plan. Though this assumption may not bias our results, it is possible that consumers may not accurately assess their future health risk nor have a comprehensive understanding of the costs and benefits of different plan options. Adapting our model to allow for heuristic decision making and/or choice under ambiguity would add to our understanding of how consumers choose health plans and, subsequently, how information provision policies impact market outcomes. Third, our model assumes that switching costs contribute to the relative utility of the default option in an additive linear manner. It is possible that switching costs occur in part due to phenomena that are unrelated to relative plan utilities (e.g. fixed re-optimization costs). Future work could investigate an empirical framework where switching costs enter as a fixed cost independent of relative plan utilities as well as in the way we have modeled them.

In addition to improving aspects of the consumer choice framework, we could delve further into the rich data we possess on several dimensions. We do not currently study the impact of health savings accounts or flexible savings accounts on plan choice behavior even though we observe detailed data on enrollment and contributions for these benefits. A careful study of consumer utilization of these savings vehicles could provide further insight into the degree of consumer choice adequacy and the determinants of plan choice. On the cost side, we could advance our analysis of moral hazard and plan selection based on private information by taking further advantage of the natural experiment that occurs in our data. The firm's menu change provides a plausibly exogenous shift in coverage for most of the population which can be used in conjunction with our detailed medical data to rigorously study how consumers respond to different marginal prices for medical services. In addition to studying variation in utilization across plan years we could study within year price elasticities in scenarios where consumers face different marginal prices due to exogenous shocks (see e.g. Kowalski (2009)). We could also investigate the extent of selection based on private information by studying the correlation between predictable ex post claims (such as pregnancy) and coverage choices. While interesting in their own right, these advances would enrich our understanding of consumer choice behavior.

Finally, this analysis assumes a simplified pricing environment that is identical to that faced by consumers within the firm we study. Future work could investigate the interaction between switching costs and adverse selection with a richer firm pricing model. Especially in the context of a health insurance exchange, where competing firms price directly to individual consumers, firms cognizant of switching costs will take them into account by pricing dynamically. Further, our analysis assumes that the menu of plan options is constant while in reality the level of coverage

may be determined endogenously. The estimates of our structural choice and cost models can be directly applied to investigate these counterfactual scenarios once we develop richer models of firm behavior. Such analyses will shed light on the important regulatory issues currently being discussed regarding health insurance exchanges and insurance market reform.

References

ABALUCK, J., AND J. GRUBER (2009): "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," NBER Working Paper No. 14759.

AKERLOF, G. (1970): "The market for 'lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84(3), 488–500.

BAJARI, P., H. HONG, AND A. KHWAJA (2006): "Moral Hazard, Adverse Selection, and Health Expenditures: A Semiparametric Analysis," NBER Working Paper No. 12445.

BUNDORF, K., J. LEVIN, AND N. MAHONEY (2008): "Pricing, Matching and Efficiency in Health Plan Choice," Stanford University.

CARDON, J., AND I. HENDEL (2001): "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *The RAND Journal of Economics*, 32(3), 408–427.

CARLIN, C., AND R. TOWN (2007): "Adverse Selection, Welfare, and Optimal Pricing of Employer Sponsored Health Plans," University of Minnesota Working Paper.

CHIAPPORI, P., AND B. SALANIE (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108, 56–78.

COHEN, A., AND L. EINAV (2007): "Estimating Risk Preferences from Deductible Choice," *The American Economic Review*, 97(3), 745–788.

CUTLER, D., A. FINKELSTEIN, AND K. MCGARRY (2008): "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," NBER Working Paper 13746.

CUTLER, D., B. LINCOLN, AND R. ZECKHAUSER (2009): "Selection Stories: Understanding Movement Across Health Plans," NBER Working Paper No. 15164.

CUTLER, D., AND S. REBER (1998): "Paying for Health Insurance: The Tradeoff Between Competition and Adverse Selection," *Quarterly Journal of Economics*, 113(2), 433–466.

DUBE, J., G. HITSCH, P. ROSSI, AND M. VITORINO (2008): "Category Pricing with State Dependent Utility," *Marketing Science*, 27(3), 417–429.

EINAV, L., A. FINKELSTEIN, AND M. CULLEN (2009): "Estimating Welfare in Insurance Markets Using Variation in Prices," Stanford University and M.I.T. working paper.

EINAV, L., A. FINKELSTEIN, AND J. LEVIN (2009): "Beyond Testing: Empirical Models of Insurance Markets," NBER Working Paper No. 15241.

EINAV, L., A. FINKELSTEIN, AND P. SCHRIMPF (2009): "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market," M.I.T. Working Paper.

ELLISON, G. (2006): "Bounded Rationality in Industrial Organization," MIT and NBER.

ENTHOVEN, A., A. GARBER, AND S. SINGER (2001): "Near-Universal Coverage Through Health Plan Competition: An Insurance Exchange Approach," *Covering America: Real Remedies for the Uninsured, Editors J. Meyer and E. Wicks*, Washington DC.

ESTEVES-SORENSON, C. (2009): "Micro-costs: Intertia in Television Viewing," Yale School of Management working paper.

FANG, H., M. KEANE, AND D. SILVERMAN (2008): "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market," *Journal of Political Economy*, 116(2), 303–350.

FARRELL, J., AND P. KLEMPERER (2007): "Coordination and Lock-In: Competition with Switching Costs and Network Effects," *Handbook of Industrial Organization Chapter 31 Editors M. Armstrong and R. Porter*, pp. 1967–2072.

FINKELSTEIN, A., AND J. POTERBA (2008): "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market," *Journal of Political Economy*, 112, 183–208.

GABAIX, X., AND D. LAIBSON (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 121(2).

GERTNER, R. (1993): "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'," *Quarterly Journal of Economics*, 108(2).

GOETTLER, R., AND K. CLAY (2009): "Tariff Choice with Consumer Learning and Switching Costs," University of Chicago Booth School of Business.

HOLT, C., AND S. LAURY (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5).

JONES, D. (2009): "Inertia and Overwithholding: Explaining the Prevalence of Income Tax Refunds," Stanford Institute for Economic Policy Research.

Kaiser Family Foundation (2009): "Employer Health Benefits: 2009 Summary of Findings," Discussion paper, The Kaiser Family Foundation and Health Research & Educational Trust.

KLEMPERER, P. (1995): "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economic Studies*, 62, 515–539.

KLING, J., S. MULLAINATHAN, E. SHAFIR, L. VERMEULEN, AND M. WROBEL (2008): "Misperception in Choosing Medicare Drug Plans," working paper.

KOWALSKI, A. (2009): "Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care," Yale Department of Economics working paper.

LUSTIG, J. (2008): "The Welfare Effects of Adverse Selection in Privatized Medicare," Boston University working paper.

MADRIAN, B., AND D. SHEA (2001): "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 141(4), 1149–1187.

MANNING, W., J. NEWHOUSE, N. DUAN, E. KEELER, AND A. LEIBOWITZ (1987): "Health Insurance and the Demand for Medical Care: Evidence From a Randomized Experiment," *The American Economic Review*, 77(3), 251–277.

METRICK, A. (1995): "A Natural Experiment in 'Jeopardy!'," *American Economic Review*, 85(1).

NEWHOUSE, J. (1993): *Lessons from the RAND Health Insurance Experiment.* Harvard University Press.

ROTHSCHILD, M., AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90, 630–649.

SHCHERBAKOV, O. (2007): "Measuring Consumer Switching Costs in the Television Industry," working paper.

STROMBOM, B., C. BUCHMUELLER, AND J. FELDSTEIN (2002): "Switching costs, price sensitiivity, and health plan choice," *Journal of Health Economics*, 21, 89–116.

SUNSTEIN, C., AND R. THALER (2008): *Nudge.* Yale University Press.

SYDNOR, J. (2008): "Sweating the Small Stuff: Risk Aversion in Home Insurance," Case Western Reserve University.

TCHERNIS, R., S. NORMAND, J. PAKES, P. GACCIONE, AND J. NEWHOUSE (2005): "Selection and Plan Switching Behavior," Indiana University working paper.

Appendix A: Cost Model Setup and Estimation

This appendix describes the details of the cost model, which is summarized earlier in section 5. To refresh, we categorize the universe of total claims into four subdivisions of claims (i) hospital and physician (ii) pharmacy (iii) mental health and (iv) physician office visit. We divide claims into these four categories so that we can accurately characterize the plan-specific mappings from total claims to out-of-pocket expenditures since each of these categories maps to out-of-pocket expenditures in a different manner. We denote this four dimensional vector of claims C_{it} and any given element of that vector $C_{d,it}$ where $d \in D$ represents one of the four categories and i denotes an individual (employee or dependent). After describing how we predict this vector of claims for a given individual, we return to the question of how we determine out-of-pocket expenditures in plan j given C_{it}. We are unaware of any previous work that reconstructs the mapping from individual-level total claims to plan-specific out of pocket expenditures at this level of detail.

Denote an individual's past year of medical diagnoses and payments by ξ_{it} and the demographics age and sex by ζ_{it}. We use the ACG software mapping, denoted A, to map these characteristics into a predicted mean level of health expenditures for the upcoming year, denoted θ:[47]

$$A : \xi \times \zeta \to \theta$$

In addition to forecasting a mean level of total expenditures, the software has an application that predicts future mean pharmacy expenditures. This mapping is analogous to A and outputs a prediction λ for future pharmacy expenses.

We use the predictions θ and λ to categorize similar groups of individuals across each of the four claims categories in vector in C_{it}. Then for each group of individuals in each claims category, we use the actual ex post realized claims for that group to estimate the ex ante distribution for each individual under the assumption that this distribution is identical for all individuals within the cell. Individuals are categorized into cells based on different metrics for each of the four elements of C:

Pharmacy:	λ_{it}
Hospital / Physician (Non-OV):	θ_{it}
Physician Office Visit:	θ_{it}
Mental Health:	$C_{MH,i,t-1}$

For pharmacy claims, individuals are grouped into cells based on the predicted future mean pharmacy claims measure output by the ACG software, λ_{it}. For the categories of hospital / physician (non office visit) and physician office visit claims individuals are grouped based on their mean predicted total future health expenses, θ_{it}. Finally, for mental health claims, individuals are grouped

---

[47]As a reminder, the ACG mapping uses medical diagnostic information as well as previous cost data to predict future expenditures for each individual. We are aware of only one previous study that incorporates diagnostic information in cost prediction for the purposes of studying plan choice (Carlin and Town (2007)).

into categories based on their mental health claims from the previous year, $C_{MH,i,t-1}$ since (i) mental health claims are very persistent over time in the data and (ii) mental health claims are uncorrelated with other health expenditures in the data. For each category we group individuals into a number of cells between 8 and 10, taking into account the tradeoff between cell size and precision. The minimum number of individuals in any cell is 73 while almost all cells have over 500 members. Thus since there are four categories of claims, each individual can belong to one of approximately 10^4 or 10,000 combination of cells.

Denote an arbitrary cell within a given category d by z. Denote the population in a given category-cell combination (d, z) by I_{dz}. Denote the empirical distribution of ex-post claims in this category for this population $\hat{G_{I_{dz}}}(\cdot)$. Then we assume that each individual in this cell has a distribution equal to a continuous fit of $\hat{G_{I_{dz}}}(\cdot)$, which we denote G_{dz}:

$$\varpi : \hat{G_{I_{dz}}}(\cdot) \to G_{dz}$$

We model this distribution continuously in order to easily incorporate correlations across d. Otherwise, it would be appropriate to use $G_{I_{dz}}$ as the distribution for each cell.

The above process generates a distribution of claims for each d and z but does not model correlation over D. It is important to model correlation over claim categories because it is likely that someone with a bad expenditure shock in one category (e.g. hospital) will have high expenses in another area (e.g. pharmacy). We model correlation at the individual level by combining marginal distributions G_{idt} $\forall$ d with empirical data on the rank correlations between pairs (d, d').[48] Here, G_{idt} is the distribution G_{dz} where $i \in I_{dz}$ at time t. Since correlations are modeled across d we pick the metric θ to group people into cells for the basis of determining correlations (we use the same cells that we use to determine group people for hospital and physician office visit claims). Denote these cells based on θ by z_θ. Then for each cell z_θ denote the empirical rank correlation between claims of type d and type d' by $\rho_{z_\theta}(d, d')$. Then, for a given individual i we determine the joint distribution of claims across D for year t, denoted $H_{it}(\cdot)$, by combining i's marginal distributions for all d at t using $\rho_{z_\theta}(d, d')$:

$$\Psi : G_{iDt} \times \rho_{z_{\theta_{it}}}(D, D') \to H_{it}$$

Here, G_{iDt} refers to the set of marginal distributions $G_{idt} \forall d \in D$ and $\rho_{z_{\theta_{it}}}(D, D')$ is the set of all pairwise correlations $\rho_{z_{\theta_{it}}}(d, d') \forall (d, d') \in D^2$. In estimation we perform Ψ by using a Gaussian copula to combine the marginal distribution with the rank correlations, a process which we describe momentarily.

The final part of the cost model maps the joint distribution H_{it} of the vector of total claims C over the four categories into a distribution of out of pocket expenditures for each plan. For each of the three plan options we construct a mapping from the vector of claims C to out of pocket expenditures OOP_j:

---

[48]It is important to use rank correlations here to properly combine these marginal distribution into a joint distribution. Linear correlation would not translate empirical correlations to this joint distribution appropriately.

$$\Omega_j : C \rightarrow OOP_j$$

This mapping takes a given draw of claims from H_{it} and converts it into the out of pocket expenditures an individual would have for those claims in plan j. This mapping accounts for plan-specific features such as the deductible, co-insurance, co-payments, and out of pocket maximums listed in table 2. I test the mapping Ω_j on the actual realizations of the claims vector C to verify that our mapping comes close to reconstructing the true mapping. Our mapping is necessarily simpler and omits things like emergency room co-payments and out of network claims. We constructed our mapping with and without these omitted categories to insure they did not lead to an incremental increase in precision. We find that our categorization of claims into the four categories in C passed through our mapping Ω_j closely approximates the true mapping from claims to out of pocket expenses. Further, we find that it is important to model all four categories described above; removing any of the four makes Ω_j less accurate. Figure 5 shows the results of our validation exercise for PPO_{250}. Actual employee out of pocket amounts are close to those predicted by Ω_j while when we add out of network expenses as a fifth category the output is very similar.

[Figure 5 about here.]

Once we have a draw of OOP_{ijt} for each i (claim draw from H_{it} passed through Ω_j) we map individual out of pocket expenditures into family out of pocket expenditures. For families with less than two members this involves adding up all the within family OOP_{ijt}. For families with more than three members there are family level restrictions on deductible paid and out of pocket maximums that we adjust for. Define a family k as a collection of individuals i_k and the set of families as K. Then for a given family out of pocket expenditures are generated:

$$\Gamma_j : OOP_{i_k, jt} \rightarrow OOP_{kjt}$$

To create the final object of interest, the family-plan-time specific distribution of out of pocket expenditures $F_{kjt}(\cdot)$, we pass the claims distributions H_{it} through Ω_j and combine families through Γ_j. $F_{kjt}(\cdot)$ is then used as an input into the choice model that represents each family's information set over future medical expenses at the time of plan choice. Eventually, we also use H_{it} to calculate total plan cost when we analyze counterfactual plan pricing based on the average cost of enrollees.

We note that the decision to do the cost model by grouping individuals into cells, rather then by specifying a more continuous form, has costs and benefits. The cost is that all individuals within a given cell for a given type of claims are treated identically. The benefit is that our method produces local cost estimates for each individual that are not impacted by the combination of functional form and the health risk of medically different individuals. Also, the method we use allows for flexible modeling across claims categories. Finally, we note that we map the empirical distribution of claims to a continuous representation because this is convenient for building in correlations in the next step. The continuous distributions we generate vert closely fit the actual empirical distribution of claims across these four categories.

Cost Model Identification and Estimation

The cost model is identified based on the two assumptions of (i) no moral hazard / selection based on private information and (ii) that individuals within the same cells for claims d have the same ex ante distribution of total claims in that category.[49] Once these assumptions are made, the model uses the detailed medical data, the Johns Hopkins predictive algorithm, and the plan-specific mappings for out of pocket expenditures to generate the the final output $F_{kjt}(\cdot)$.

Once we group individuals into cells for each of the four claims categories, there are two statistical components to estimation. First, we need to generate the continuous marginal distribution of claims for each cell z in claim category d, G_{dz}. To do this, we fit the empirical distribution of claims $G_{I_{dz}}$ to a Weibull distribution with a mass of values at 0. We use the Weibull distribution instead of the log normal distribution, which is traditionally used to model medical expenditures, because we find that the log normal distribution over predicts large claims in the data while the Weibull does not. For each d and z the claims greater than zero are estimated with a maximum likelihood fit to the Weibull distribution:

$$\max_{(\alpha_{dz},\beta_{dz})} \Pi_{i \in I_{dz}} \frac{\beta_{dz}}{\alpha_{dz}} \left(\frac{c_{id}}{\alpha_{dz}}\right)^{\beta_{dz}-1} e^{-\left(\frac{c_{id}}{\alpha_{dz}}\right)^{\beta_{dz}}}$$

Here, $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ are the shape and scale parameters that characterize the Weibull distribution. Denoting this distribution $W(\hat{\alpha_{dz}}, \hat{\beta_{dz}})$ the estimated distribution $\hat{G_{dz}}$ is formed by combining this with the estimated mass at zero claims, which is the empirical likelihood:

$$\hat{G_{dz}}(c) = \begin{cases} G_{I_{dz}}(0) & \text{if } c = 0 \\ G_{I_{dz}}(0) + \frac{W(\hat{\alpha_{dz}},\hat{\beta_{dz}})(c)}{1 - G_{I_{dz}}(0)} & \text{if } c > 0 \end{cases}$$

Again, we use the notation $\hat{G_{iDt}}$ to represent the set of marginal distributions for i over the categories d: the distribution for each d depends on the cell z an individual i is in at t. We combine the distributions $\hat{G_{iDt}}$ for a given i and t into the joint distribution H_{it} using a Gaussian copula method for the mapping Ψ. Intuitively, this amounts to assuming a parametric form for correlation across $\hat{G_{iDt}}$ equivalent to that from a standard normal distribution with correlations equal to empirical rank correlations $\rho_{z_{\theta_{it}}}(D, D')$ described in the previous section. Let $\Phi^i_{1|2|3|4}$ denote the standard multivariate normal distribution with pairwise correlations $\rho_{z_{\theta_{it}}}(D, D')$ for all pairings of the four claims categories D. Then an individual's joint distribution of non-zero claims is:

$$\hat{H_{i,t}}(\cdot) = \Phi_{1|2|3|4}(\Phi_1^{-1}(\hat{G_{id_1t}}), \Phi_2^{-1}(\hat{G_{id_2t}}), \Phi_3^{-1}(\hat{G_{id_3t}}), \Phi_4^{-1}(\hat{G_{id_4t}}))))$$

Above, Φ_d is the standard marginal normal distribution for each d. $\hat{H}_{i,t}$ is the joint distribution of claims across the four claims categories for each individual in each time period. After this is estimated, we determine our final object of interest $F_{kjt}(\cdot)$ by simulating K multivariate draws

---

[49]Robustness checks on both of these assumptions are an important path for future work.

from $\hat{H}_{i,t}$ for each i and t, and passing these values through the plan-specific total claims to out of pocket mapping Ω_j and the individual to family out of pocket mapping Γ_j. The simulated $F_{kjt}(\cdot)$ for each k, j, and t is then used as an input into estimation of the choice model as described in section 5.
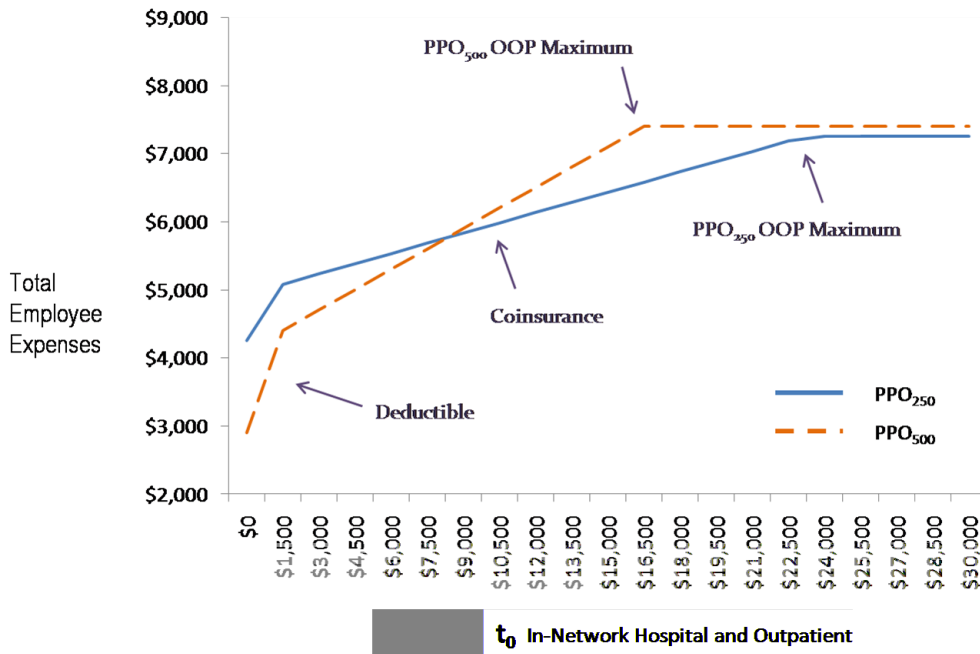
Figure 1: This figure describes the relationship between total medical expenses (plan plus employee) and employee expenses in year t_0 for PPO_{250} and PPO_{500}. This mapping depends on employee premiums, deductible, coinsurance, and out of pocket maximum. This chart applies to low income families (premiums vary by number of dependents covered and income tier, so there are similar charts for all 20 combinations). Healthy individuals should choose PPO_{500} and sick individuals should choose PPO_{250}. Premiums are treated as pre-tax expenditures while medical expenses are treated as post-tax.
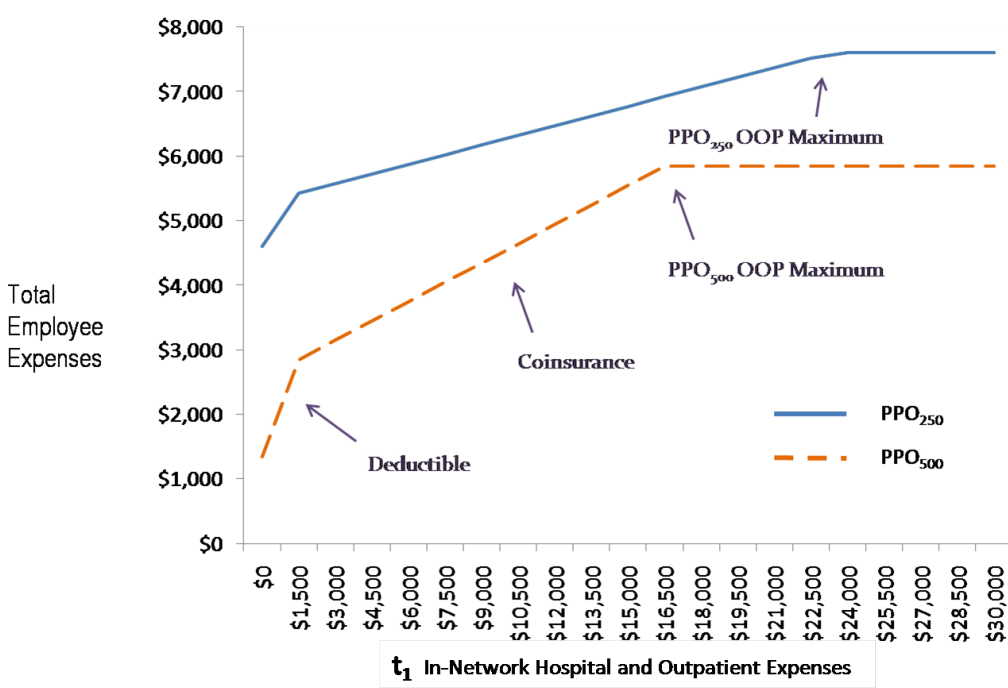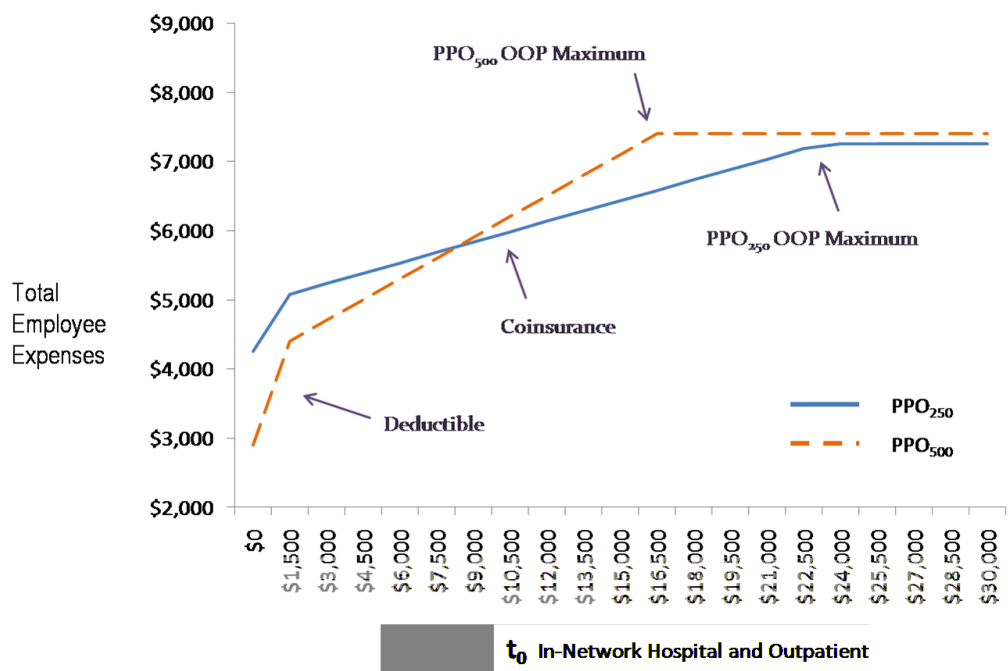
Figure 2: The top panel repeats Figure 1 for comparison purposes and describes employee expenses as a function of total hospital and physician claims in t_0 for a low income family. The bottom panel presents the analogous chart for time t_1 when premiums changed significantly. This can be seen by the change in the vertical intercepts. At time t_0 healthier employees were better off in PPO_{500} and sicker employees were better off in PPO_{250}. For this group, at time t_1 *all* employees should choose PPO_{500} regardless of their total claim levels, i.e. PPO_{250} is dominated by PPO_{500}. Despite this, many employees who chose PPO_{250} in t_0 continue to do so at t_1, indicative of high switching costs.

49
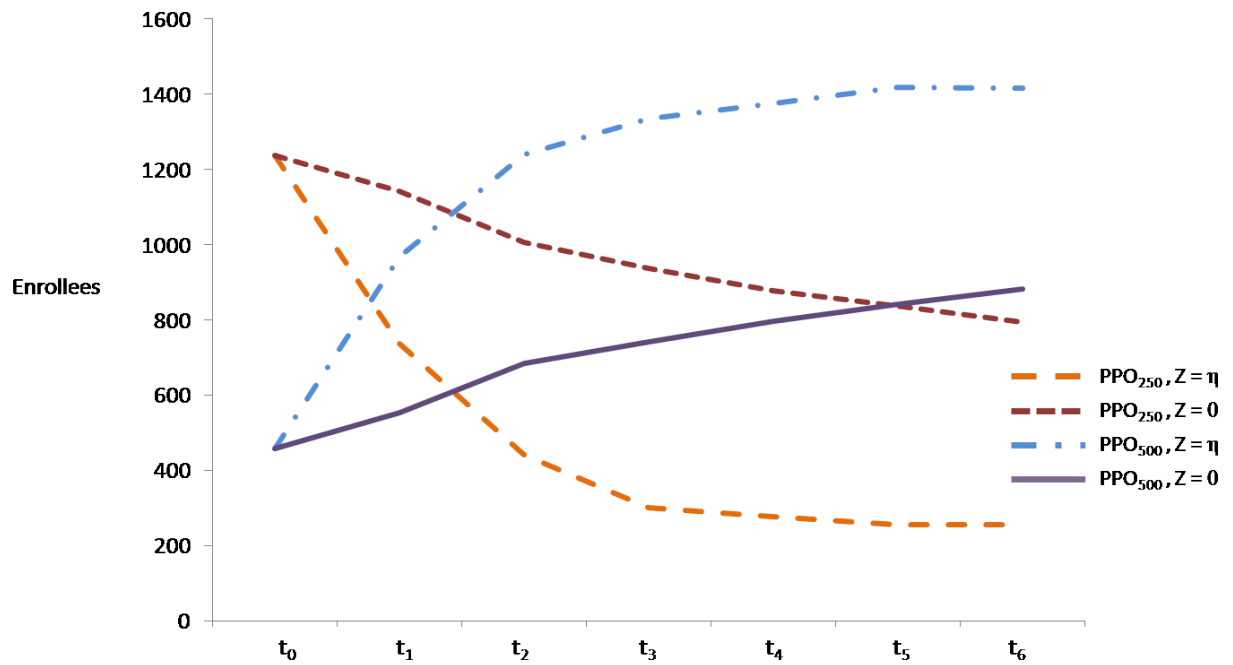
Figure 3: This figure presents the time path of choices for PPO_{250} and PPO_{500} with and without the policy intervention to reduce switching costs. In full equilibrium, the impact of the policy intervention on the market share of PPO_{250} relative to PPO_{500} is noticeable. In the benchmark case where there are significant switching costs over the six year period the market share of PPO_{250} declines from 1237 to 793 while that of PPO_{500} increases from 458 to 882. After the policy intervention reduces switching costs to zero, PPO_{250} enrollment declines all the way to 255 after six years while PPO_{500} enrollment increases to 1416.
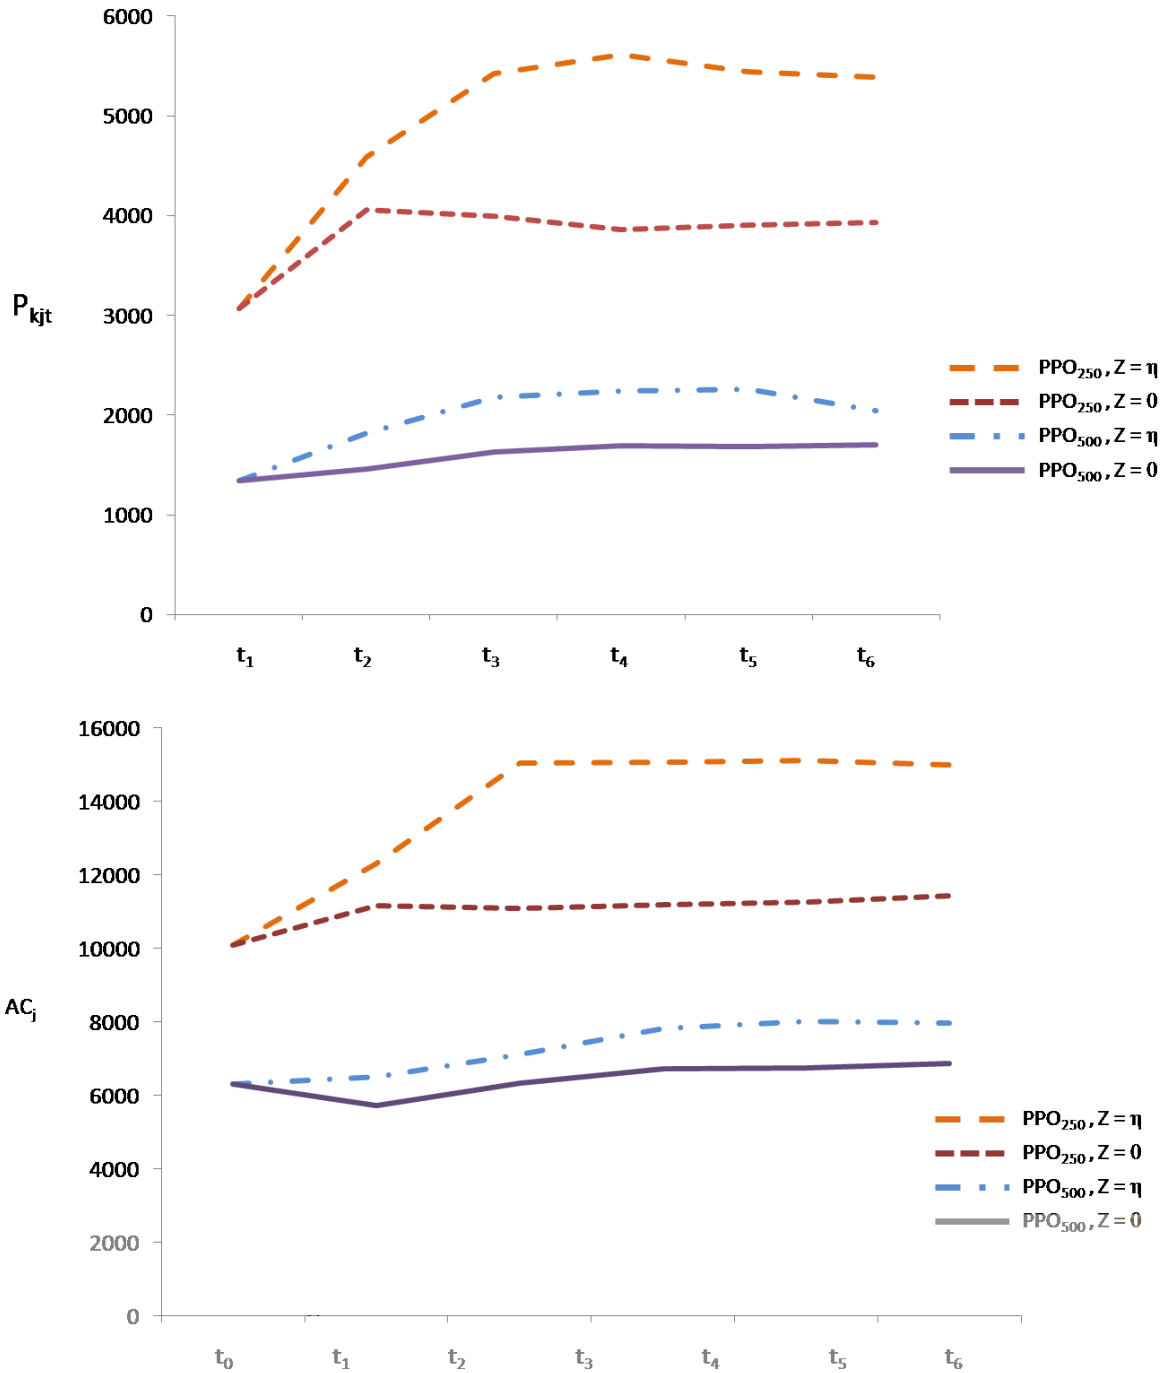
Figure 4: This figure presents the time path of average costs and prices for single individuals enrolled in PPO_{250} and PPO_{500} with and without the policy intervention. Prices are given for individuals in the middle income tier. When switching costs are reduced to zero, the employee premium of PPO_{250} increases from \$3,071 at t_1 to \$5,390 while in the benchmark model with switching costs the long run price is \$3,929. The long run price of $PP0_{500}$ without switching costs is \$2,044 and with is \$1,733.
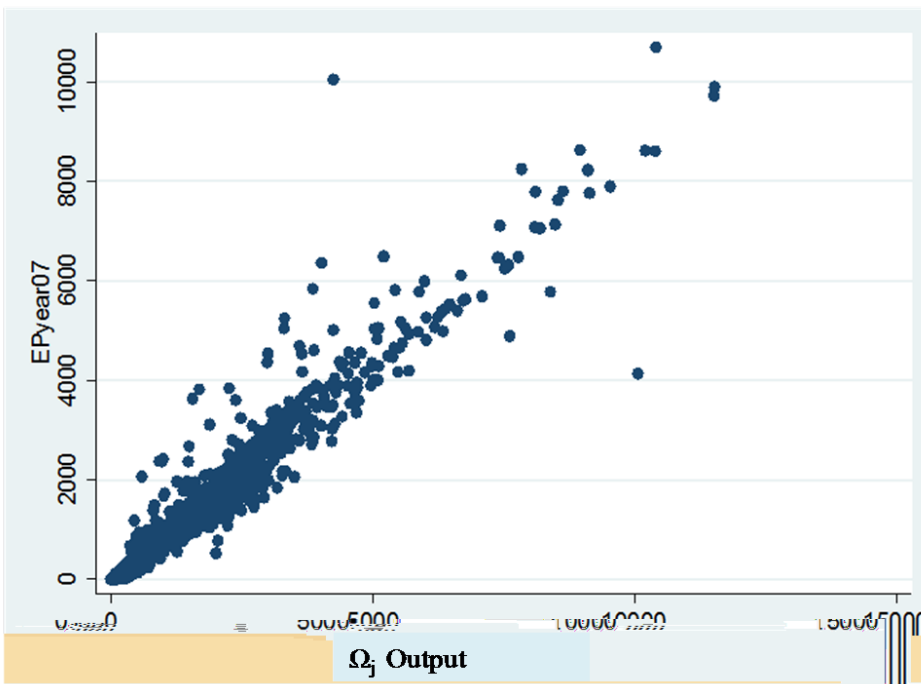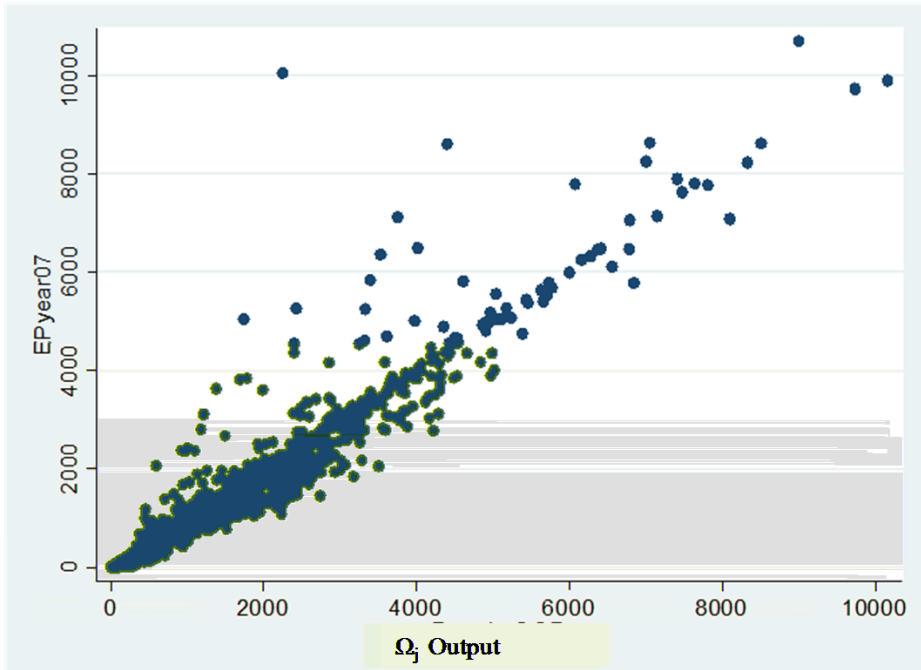
Figure 5: This figure validates the mapping Ω_j that translates the vector of total claims C into plan specific out of pocket expenditures. The two charts show Ω for PPO_{250}. The top chart is the mapping actually used where claims are categorized into four categories (i) hospital and outpatient (ii) pharmacy (iii) mental health and (iv) physician office visit. Ideally, we want all points to be on the 45 degree line where the actual employee paid out of pocket equals the model predicted out of pocket. The plot is condensed around the 45 degree line so we believe this our mapping is close to the true mapping. The bottom figure adds out of network expenses to the mapping as a fifth category and does not materially improve upon the mapping used.

	All Employees	PPO Ever 04-09	Final Sample
EMPLOYEES	14,248	6,398	2,022
GENDER (MALE %)	47.4%	45.9%	48.5%
MEAN AGE (MEDIAN)	39.9	39.9	46
	(37)	(37)	(46)
INCOME			
Tier 1	31.3%	31.7%	20.3%
Tier 2	36.6%	39.4%	41.4%
Tier 3	17.3%	18.5%	23.9%
Tier 4	6.5%	5.6%	7.5%
Tier 5	8.3%	4.8%	6.9%
FAMILY SIZE			
1	59.9 %	57.1 %	44.5 %
2	15.5 %	18.4 %	21.2 %
3	10.4 %	10.7 %	13.9 %
4+	14.2 %	13.8 %	27.9 %
STAFF GROUPING			
MANAGER	25.7%	24.3%	34.3%
WHITE-COLLAR	46.1%	47.5%	43.1%
BLUE-COLLAR	28.3%	27.9%	21.7%
EMPLOYMENT CHARACTERISTICS			
QUANTITATIVE MANAGER	16.6%	13.1%	19.2%
JOB TENURE (MEDIAN YEARS)	4.6	3.8	7.8

Table 1: The first column describes demographics for the entire sample whether or not they ever enroll in insurance with the firm. A higher numbered income tier implies higher income. The second column summarizes demographics for the sample of individuals who ever enroll in a PPO option (people who ever appear in the claims data). The third column describes our final estimation sample which includes those employees who (i) are enrolled in PPO_{-1} at t_{-1} and (ii) remain enrolled in any plan at the firm through at least t_2. The final estimation sample is slightly older, richer, and more family oriented then the full sample.

	PPO_{-1}	PPO_{250}	PPO_{500}	PPO_{1200}
IND. DEDUCTIBLE (FAMILY)	250*	250	500	1200
	(750)	(750)	(1500)	(2400)
CO-INSURANCE	10%	10%	20%	20%
PHY. VISIT CO-PAY	20	25	25	NA
ER CO-PAY	100	100	100	NA
MENTAL HEALTH CI	50%	50%	50%	50%
PHARMACY CO-PAY	5/25/45**	5/25/45**	5/25/45**	NA
	(10/45/65)	(10/50/75)	(10/50/75)	NA
IND. OOP MAX (FAMILY)***				
Income tier 1	2000	1000	1500	2000
	(6000)	(3000)	(4500)	(6000)
Income tier 2/3	2000	2000	3000	4000
	(6000)	(5000)	(7000)	(8000)
Income tier 4/5	2000	3000	4000	5000
	(6000)	(8000)	(9000)	(10000)

* PPO_{-1} has inpatient deductible of 150 per admission

** Prescription max of 1500 per person

*** Office visit and pharmacy claims only apply to OOP max for PPO_{1200}

Table 2: This table describes the financial characteristics for each PPO option that determine how much an individual pays for medical expenses out of pocket. For most medical expenses, an individual pays the full amount until he reaches the yearly plan deductible, after which he pays the coinsurance rate for all further medical expenses. Once an individual spends the out of pocket maximum, he pays no further general medical expenses. Pharmacy products and physician office visits only apply to the deductible and coinsurance for PPO_{1200}; all other plans have fixed co-payments for these services. Mental health services apply to all plan deductibles (but not OOP max) and have 50% coinsurance post deductible. Out of pocket maximums vary with income, presumably for equity considerations. This chart does not include out of network plan characteristics. Overall, out of network expenses account for only 2% of total expenses.

Income - Dependents	Single t_0	Single t_1	Family t_0	Family t_1
PPO_{250}				
Tier 1	1308	1512	**4260**	**4596**
Tier 2	1476	1800	4812	5532
Tier 3	1956	2184	6360	6780
Tier 4	2556	2664	8328	8340
Tier 5	2964	3384	9660	10680
PPO_{500}				
Tier 1	888	384	**2904**	**1344**
Tier 2	1056	672	3456	2352
Tier 3	1536	1056	5004	2808
Tier 4	2136	1536	6972	4368
Tier 5	2544	2256	8304	6708
PPO_{1200}				
Tier 1	132	120	468	420
Tier 2	300	240	1020	840
Tier 3	780	600	2568	2100
Tier 4	1380	960	4536	3360
Tier 5	1788	1464	5868	4200

Table 3: This table describes employee premium contributions for t_0 and t_1 (the first two periods after the plan menu change). There are significant price changes from one year to the next from the perspective of employees. For example, for a family with low income, the premium for PPO_{500} dropped $1,560 (from $2,904 to $1,344) from t_0 to t_1 while the premium for PPO_{250} increased by $336 (from $4,260 to $4,596) over that time period.

	t_{-1}	t_0	t_1
PPO_{250}	-	2,199 (25%)	1,937 (21%)
PPO_{500}	-	998 (11%)	1,544 (18%)
PPO_{1200}	-	876 (10%)	824 (9%)
HMO_1	2,094 (25%)	2,050 (23%)	2,031 (22%)
HMO_2	701 (8%)	1,273 (14%)	1,181 (13%)
PPO_{-1}	3,264 (39%)	-	-
HMO_3	668 (8%)	-	-
HMO_4	493 (6%)	-	-
Waive	1,207 (14%)	1,447 (16%)	1,521 (17%)

Table 4: This table describes the choices of all employees at the firm from year t_{-1} (right before the menu change) to year t_1 (one year after the menu change). Among the PPO options after the menu change, PPO_{250} has the largest market share in years t_0 and t_1. The number of employees enrolled in both the PPO and HMO nests of plans is close to constant over time after the menu change. Approximately 15% of employees waive coverage in each year.

t_0 plan / t_{-1} plan	PPO_{-1}	All HMO
PPO_{250}	1,710	194
PPO_{500}	570	118
PPO_{1200}	392	147
HMO_1	49	1,703
HMO_2	36	943

t_1 plan / t_0 plan	PPO_{250}	PPO_{500}	PPO_{1200}	HMO_1	HMO_2
PPO_{250}	1,732	14	14	21	8
PPO_{500}	129	774	112	31	31
PPO_{1200}	17	11	577	12	15
HMO_1	10	7	8	1,694	
HMO					

	New Enrollee t_{-1}	New Enrollee t_0	New Enrollee t_1
t_0			
PPO_{250}	259 **(25%)**	287 **(21%)**	-
PPO_{500}	205 **(19%)**	306 **(23%)**	-
PPO_{1200}	155 (15%)	236 (17%)	-
HMO_1	238 (23%)	278 (20%)	-
HMO_2	199 (18%)	270 (19%)	-
t_1			
PPO_{250}	182 **(23%)**	253 **(20%)**	142 **(11%)**
PPO_{500}	201 **(26%)**	324 **(26%)**	562 **(43%)**
PPO_{1200}	95 (12%)	194 (15%)	188 (14%)
HMO_1	171 (22%)	257 (20%)	262 (20%)
HMO_2	135 (17%)	239 (19%)	151 (12%)
Demographics			
Mean Age	33	33	32
Median Age	31	31	31
Income Tier 1 (t_1)	48%	50%	47%
Income Tier 2	33%	31%	32%
Income Tier 3	10%	10%	12%
Income Tier 4	5%	4%	4%
Income Tier 5	4%	5%	5%

Table 6: This table describes the choice behavior of new entrants to the firm over several consecutive years. We hypothesize that, if there are no switching costs, the t_1 choices of employees who entered the firm at t_0 should be identical to the t_1 choices of employees who entered the firm at t_1. This table shows that this hypothesis does not hold; the t_1 choices of employees who enter at t_0 reflect both t_1 prices and t_0 choices while the t_1 choices of new employees at t_1 reflect just t_1 prices. New employees at t_0 do not adjust to the significant price change from t_0 to t_1 while new employees' choices do reflect these price changes. This is true even though these populations have almost identical demographic profiles.

	Dominated PPO_{250}	Dominated PPO_{250} Switchers
N	708	73
PPO_{500}, t_1	-	50 (69%)
PPO_{1200}, t_1	-	6 (8%)
Other	-	17 (23%)
FSA 2008 Enrollment %	**29%**	**53%**
Dental Switch	**3.6%**	**9.5%**
Mean Income Tier	2.4	2.2
Quantitative Manager	18%	11%
Mean Age	46.8	40.8
Single %	43%	57%
Mean Family Size	2.1	2.02

Table 7: This table profiles the choices and demographics of the employees enrolled in PPO_{250} at t_0 who (i) continue to enroll in a firm plan in t_1 and (ii) have PPO_{250} become dominated for them at t_1. The majority of these employees (635 out of 708) remain in PPO_{250} even after it becomes dominated by PPO_{500}. People who do switch are more likely to exhibit a pattern of active choice behavior in general as evidence by their higher FSA enrollments and level of dental plan switching. Apart from this, these populations are similar though switchers in this group are slightly younger and more likely to be single.

	N	Mean Fam Size	Mean	25th pct	Median	75th pct
PPO_{-1}	2022	2.24	13,331	1,257	4,916	13,022
$PPO_{250}\ t_0$	1,328	2.18	16,976	2,041	6,628	16,135
$PPO_{500}\ t_0$	338	2.20	6,151	554	2,244	6,989
$PPO_{1200}\ t_0$	280	2.53	6,742	658	2,958	8,073
$PPO_{250}\ t_1$	1,244	2.19	17,270	2,041	6,651	16,707
$PPO_{500}\ t_1$	461	2.19	7,759	708	2,659	8,588
$PPO_{1200}\ t_1$	232	2.57	6,008	589	2,815	7,191

Table 8: This table shows the extent of adverse selection across PPO options after the t_0 menu change. The numbers in this table represent t_{-1} total claims in dollars. All individuals in this sample were enrolled in PPO_{-1} in t_{-1} and continue to be enrolled in some plan at the firm for the following two years. It is clear that for any distributional feature, more expensive and risky families choose $PP0_{250}$ in both t_0 and t_1. When prices change significantly from t_0 to t_1 the pattern of selection does not change. The lack of incremental selection can be attributed to switching costs that stop people from moving optimally across plans at t_1.

Control				Treatment		
	t_{-1}	t_0	%	t_{-1}	t_0	%
Aggregate Expenses						
25th Pctile	$2,371	$2,591	9%	$808	$994	23%
Median	$6,985	$7,564	8%	$2,852	$3,130	10%
75th Pctile	$16,827	$17,909	7%	$8,020	$9,442	17%
Mean	$17,531	$17,156	-3%	$6,816	$8,493	21%
Count	1344			642		

Table 9: This table analyzes the medical consumption behavior of consumers who select less comprehensive coverage at t_0 and those who select more comprehensive coverage. Those who select more comprehensive insurance (PPO_{250}) face similar incentives to those they faced in PPO_{-1} before the menu change (control group) while those who enroll in PPO_{500} or PPO_{1200} pay higher marginal prices for medical care (treatment group). This table shows that the expenses of those in the treatment group increase by more than those of the control group which is the opposite to the effect that moral hazard or time-varying selection based on private information would predict.

	Med^{250}_{-1}	Ratio^{C}	Ratio^{T}	$\Delta\mathrm{Ratio}$	MH/PI
Diagnostic Category					
Benign / Uncertain Neoplasm	$297	5.7%	26.8%	-21.11%	N
Diabetes	$ 290	-8.2%	22.3%	-30.6%	N
Ears, Nose & Throat	171$	-1.1%	20%	-21.17%	N
Eyes	$170	16.5%	28.5%	-12.1%	M
Gastrointestinal	$447	-13%	-52%	39%	Y
Genital System	$186	-5.4%	30.5%	-35.9%	N
Heart	$272	1.1%	-34.2%	35.3%	Y
Hematological	$159	-25.8%	80.7%	-106.7%	N
Infectious	$129	8.5%	51.5%	-43%	N
Injury / Poisoning	$714	-8.4%	-9.45%	1.1%	M
Lung	$130	10.8%	6.1%	4.6%	M
Malignant Neoplasm	$1,777	-33.7%	16.1%	-49.9%	N
Mental	$1,233	-10.3%	-26.9%	16.6%	M
Musculoskeletal	$860	2.1%	-7.3%	9.5%	M
Nutritional / Metabolic	$170	1.2%	35.5%	-34.3%	N
Pregnancy	$4,246	12%	-73%	85%	Y
Screening	$339	23.3%	19.3%	4%	N
Skin	$171	6.4%	10.8%	-4.4%	M
Symptoms / Signs	$468	2.6%	-2.7%	5.3%	M
Urinary System	$128	-3.9%	31.7%	-35.6%	N

Table 10: This table presents analogous results to Table 9 for specific categories of medical expenditures. The table presents an analysis of claims in each category conditional on claims being larger than zero. There is little difference from year to year of the total claim count in each category for each plan. The column Ratio^{C} presents the difference in difference ratio of the median claims > 0 in each category from t_{-1} to t_0 while Ratio^{T} presents this statistic for people in the treatment group (who enroll in less comprehensive coverage in t_0). $\Delta\mathrm{Ratio}$ presents the difference in these ratios, which will be > 0 when there is moral hazard or selection based on private information. The final column indicates whether there is evidence for either in each diagnostic category where Y stands for yes N stands for no and M implies the difference is close to zero.

	Median Regression	Mean Regression
N	5362	5362
α	0.35* (0.0126)	0.42* (0.0102)
β	-0.021* (0.0073)	-0.017* (0.0061)
R^2	0.1651	0.2073

* Denotes statistical significance at a 5% level

Table 11: This table describes the results of the difference in difference regression comparing the changes in utilization of the control group and treatment group across t_{-1} and t_0. α is the level of next year's claims predicted by this year's claims for a given individual within a given medical diagnostic category and β is the reduction in next year's claims that results solely from being enrolled in one of the less comprehensive options at t_0. In both specifications the combined effect of moral hazard and selection based on private information is $\frac{\beta}{\alpha}$, which is less than 5% in either case.

	PPO$_{-1}$	PPO$_{250}$	PPO$_{500}$	PPO$_{1200}$
Pharmacy				
Mean	$973	$1420	$586	$388
Median	$81	$246	$72	$22
Mental Health (>0)				
Mean	$2401	$2228	$1744	$2134
Median	$1260	$1211	$1243	$924
Hospital / Physician				
Mean	$4588	$5772	$2537	$2722
Median	$428	$717	$255	$366
Physician OV				
Mean	$461	$571	$381	$223
Median	$278	$356	$226	$120

Table 12: This table describes the mean and median of claims across the four modeled claims categories (i) Pharmacy (ii) Mental Health (iii) Hospital / Physician and (iv) Office Visit.

Parameter	Normal γ	Log-Normal γ
Switching Cost Individual, η_f	1570 (132)	1991 (165)
Switching Cost Family, η_s	2507 (160)	2637 (201)
Risk Aversion Mean - Intercept , μ	$4.73*10^{-4}$ * $(4.4*10^{-5})$	-8.61 (0.23)
Risk Aversion Mean - Income Slope , β	$7.71*10^{-5}$ $(9.0*10^{-6})$	0.24 (0.02)
Risk Aversion Std. Deviation , σ_γ	$3.33*10^{-4}$ $(3.6*10^{-5})$	1.22 (0.10)
PPO_{1200}-Mean Individual	-4993 (190)	-3613 (175)
PPO_{1200}-Std. Error Individual	1797 (151)	1310 (140)
PPO_{1200}-Mean Family	-5148 (201)	-5519 (283)
PPO_{1200}-Std. Error Family	2148 (130)	2256 (155)
Single High Cost Intercept PPO_{500}	-758 (279)	-917 (333)
Single High Cost Intercept PPO_{1200}	-2212 (692)	-1880 (745)
Family High Cost Intercept PPO_{500}	-1655 (544)	-1772 (620)
Family High Cost Intercept PPO_{1200}	-3506 (1224)	-3373 (1267)
ϵ_{500}	356 (62)	329 (88)
ϵ_{1200}	1002 (188)	554 (120)

* We truncate 4% of the normal distribution of γ at 0 since this parameter is > 0 in the CARA model.

Table 13: This table presents the results from the structural choice model. The first specification presented assumes normal risk aversion heterogeneity while the second assumes log normal heterogeneity. All non-risk aversion coefficients are in dollar units. Standard errors, in parentheses, are low for all parameters reflecting the high precision of the estimates. For the risk aversion estimates in the log normal model, μ represents the the first parameter in the log normal distribution while σ_γ represents the second (these are not the mean and variances of this distribution). The key finding is that switching costs are large and stable across both of these specifications.

	Absolute Risk Aversion	Interpretation
Normal Heterogeneity		
Mean / Median Individual	$6.94 * 10^{-4}$	93.6
25th percentile	$4.69 * 10^{-4}$	94.0
75th percentile	$9.19 * 10^{-4}$	91.5
90th percentile	$1.12 * 10^{-3}$	89.8
95th percentile	$1.24 * 10^{-3}$	88.9
99th percentile	$1.47 * 10^{-3}$	86.6
Log normal Heterogeneity		
Mean	$7.88 * 10^{-4}$	92.6
25th percentile	$1.64 * 10^{-4}$	97.1
Median	$3.74 * 10^{-4}$	95.2
75th percentile	$8.52 * 10^{-4}$	92.0
90th percentile	$1.79 * 10^{-3}$	84.1
95th percentile	$2.79 * 10^{-3}$	78.1
99th percentile	$6.40 * 10^{-3}$	60.5
Comparable Estimates		
Cohen-Einav Benchmark Mean	$3.1 * 10^{-3}$	76.5
Cohen-Einav Benchmark Median	$3.4 * 10^{-5}$	99.7
Gertner (1993)	$3.1 * 10^{-4}$	97.0
Metrick (1995)	$6.6 * 10^{-5}$	99.3
Holt and Laury (2002)	$3.2 * 10^{-2}$	21.0
Sydnor (2006)	$2.0 * 10^{-3}$	83.3

Table 14: This table examines the estimated risk preferences. The interpretation column is the value X that would make someone indifferent about accepting a 50-50 gamble where you win \$100 and lose X versus a status quo where nothing happens. Our estimates are similar under both specifications with the exception that the log normal model predicts a fatter tail with higher risk aversion. These estimates are in the middle of the range found in the literature and show a moderate degree of risk aversion.

	$Z = 0$ (Benchmark)	$Z = \frac{\eta}{2}$	$Z = \eta$ (No SC)
t_1 **Choices**			
PPO_{250}	1,221	1,138	852
PPO_{500}	504	594	910
PPO_{1200}	194	185	155
t_2 **Choices**			
PPO_{250}	1,160	1,037	797
PPO_{500}	573	702	994
PPO_{1200}	185	179	126
t_1 **Family Average Cost**			
PPO_{250}	26,794	28,856	30,450
PPO_{500}	17,195	17,271	19,106
PPO_{1200}	15,838	16,518	17,447
t_2 **Family Average Cost**			
PPO_{250}	27,796	31,154	31,265
PPO_{500}	17,563	18,415	20,496
PPO_{1200}	16,922	17,681	16,579

Table 15: This table presents the results of the partial equilibrium policy simulations. There are three simulations presented (i) the benchmark case with full switching costs (ii) the case when switching cost are reduced by half and (iii) the case where switching costs are completely removed. The removal of switching costs helps consumer make better decisions as more consumers adjust to the price change in favor of PPO_{500} in both cases where switching costs are reduced. At t_2 enrollment in PPO_{500} increases by 72% to 994 when switching costs are reduced to zero relative to the benchmark case. Finally, average costs for families (employees covering a spouse and at least one child) in PPO_{250} (as well as the other plans) increase when switching costs are reduced implying that additional adverse selection could be a concern when prices adjust in equilibrium.

	t_1	t_2
Mean Δ CEQ		
Population	192	215
Switchers Only	367	394
Mean Welfare Change: % Total Premiums		
Mean Employee Premium (MEP)	2,233	2,078
Welfare Change Population	8.6%	10.3%
Welfare Change Switchers	16.4%	19.0%
Mean Welfare Change: % Total Emp. Spending		
Mean Total Emp. Spending	4,305	4,375
Welfare Change Population	4.5%	5.1%
Welfare Change Switchers	8.5%	9.0%
Mean Welfare Change: % $\|CEQ\|$		
Mean Total $\|CEQ\|$	6,251	6,381
Welfare Change Population	3.2%	3.6%
Welfare Change Switchers	5.8%	6.3%

Table 16: This table presents the welfare results of the partial equilibrium policy simulations. We present the dollar change in certainty equivalents and welfare resulting from the policy intervention that reduces switching costs to 0 from η . We present three alternative welfare metrics that use a certainty equivalent based approach. These metrics divide the change in certainty equivalent from the policy intervention by (i) total employee premiums (ii) total employee spending and (iii) the absolute value of the certainty equivalent. Note that since all figures are losses the certainty equivalent absolute value is larger than the total spending figure.

	t_1	t_2	t_4	t_6
Mean Δ CEQ				
Population	$170	$117	-$120	-$132
Switcher Pop. %	30%	53%	52%	49%
Switchers Only	$567	$580	$ 360	$289
Non-Switchers Only	-$1	-$409	-$569	-$592
Healthy Pop. %	83%	84%	84%	84%
Healthy Only	$165	$130	-$137	-$123
Unhealthy Only	$195	$48	-$30	-$193
Single Pop. %	56%	55%	55%	55%
Single	$337	$268	$19	-$62
w/ Dependents	-$25	-$59	-$327	-$217
Mean Welfare Change: % Total Premiums				
Mean Employee Premium (MEP)	2,133	2,326	2,342	2,218
Welfare Change Population	7.9%	5.0%	-5.1%	-5.9%
Welfare Change Switchers	26.6%	24.9%	15.4%	13.0%
Welfare Change Non-Switchers	0%	-17.6%	-24.3%	-26.7%
Mean Welfare Change: % Total Emp. Spending				
Mean Total Emp. Spending	4,253	4,678	4,739	4,646
Welfare Change Population	4.0%	2.5%	-2.5%	-2.8%
Welfare Change Switchers	13.3%	12.4%	7.6%	6.2%
Welfare Change Non-Switchers	0%	-8.7%	-11.9%	-12.7%
Mean Welfare Change: % $\|CEQ\|$				
Mean Total $\|CEQ\|$	6,251	6,381	6,552	6,540
Welfare Change Population	2.7%	1.8%	-1.8%	-2.0%
Welfare Change Switchers	9.2%	8.9%	5.4%	4.4%
Welfare Change Non-Switchers	0%	-6.5%	-8.2%	-8.9%

Table 17: This table presents the welfare results of the full equilibrium policy simulations. We present the change in certainty equivalents and welfare resulting from the policy intervention that reduces switching costs to 0 from η. We study the distributional effects based on the three categorizations (i) 'switchers', or people who are in a different plan at time t under the policy intervention than without it (ii) an indicator of whether or not the family has high health costs relative to its coverage tier and (iii) whether an employee is single or covers dependents. We present the same welfare metrics as in table 16 taking the ratio of the change in certainty equivalent with respect to (i) total employee premiums (ii) total employee spending and (iii) the absolute value of the certainty equivalent.