NBER WORKING PAPER SERIES

ADVERSE SELECTION AND SWITCHING COSTS IN HEALTH INSURANCE MARKETS: WHEN NUDGING HURTS

Benjamin R. Handel

Working Paper 17459
http://www.nber.org/papers/w17459

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2011

I especially thank my dissertation committee chairs Igal Hendel and Michael Whinston for their guidance on this project. My third committee member, David Dranove, provided invaluable advice throughout the research process. I would like to thank conference discussants Gautam Gowrisankaran, Jon Gruber, Amanda Kowalski, and Bob Town for their advice and effort. This work has benefited particularly from the extensive comments of Severin Borenstein, Leemore Dafny, J.P. Dube, Liran Einav, Amy Finkelstein, Craig Garthwaite, Ron Goettler, Kate Ho, Mitch Hoffman, Kei Kawai, Jon Kolstad, Jon Levin, Neale Mahoney, Kanishka Misra, Aviv Nevo, Mallesh Pai, Ariel Pakes, Rob Porter, James Roberts, Bill Rogerson, Steve Tadelis, and Glen Weyl. I have received invaluable advice from numerous others including my colleagues at Berkeley and Northwestern as well as seminar participants at the 2011 AEA Meetings, ASHE-Cornell, the Bates White Antitrust Conference, Berkeley, Booth School of Business, CalTech, Columbia, the Cowles Structural Microeconomics Conference, EUI, Haas School of Business, Harvard, Harvard Business School, Harvard Kennedy School, the HEC Montreal Health-IO Conference, the Milton Friedman Institute Health Economics Conference, M.I.T., Microsoft Research, Northwestern, Olin School of Business, Princeton, Sloan School of Management, Stanford Business School, Toulouse School of Economics, UCSD, UC-Davis, University of Chicago, University of Michigan, University of Warwick, Yale and Yale SOM. I gratefully acknowledge funding support from the Center for the Study of Industrial Organization (CSIO) at Northwestern and the Robert Wood Johnson Foundation. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts
Benjamin R. Handel
NBER Working Paper No. 17459
September 2011
JEL No. D81,D82,D83,G22,I11,I18

ABSTRACT

This paper investigates consumer switching costs in the context of health insurance markets, where adverse selection is a potential concern. Though previous work has studied these phenomena in isolation, they interact in a way that directly impacts market outcomes and consumer welfare. Our identification strategy leverages a unique natural experiment that occurred at a large firm where we also observe individual-level panel data on health insurance choices and medical claims. We present descriptive results to show that (i) switching costs are large and (ii) adverse selection is present. To formalize this analysis we develop and estimate a choice model that jointly quantifies switching costs, risk preferences, and ex ante health risk. We use these estimates to study the welfare impact of an information provision policy that nudges consumers toward better decisions by reducing switching costs. This policy increases welfare in a naive setting where insurance plan prices are held fixed. However, when insurance prices change endogenously to reflect updated enrollee risk pools, the same policy substantially exacerbates adverse selection and reduces consumer welfare, doubling the existing welfare loss from adverse selection.

Benjamin R. Handel
Department of Economics
University of California, Berkeley
508-1 Evans Hall #3880
Berkeley, CA 94720
and NBER
handel@berkeley.edu

1 Introduction

A number of potential impediments stand in the way of efficient health insurance markets. The most noted of these is adverse selection, first studied by Akerlof (1970) and Rothschild and Stiglitz (1976). In insurance markets, prices reflect the expected risk (costs) of the insured pool. Whether the reason is price regulation or private information, when insurers cannot price all risk characteristics riskier consumers choose more comprehensive health plans. This causes the equilibrium prices of these plans to rise and healthier enrollees to select less comprehensive coverage than they would otherwise prefer.

A second less studied, but potentially important, impediment is poor health plan choice by consumers. A collection of research summarized by Thaler and Sunstein (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. 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. In the recently passed Affordable Care Act (ACA), policymakers emphasized clear and simple standardized insurance benefit descriptions as one way to improve consumer choices from plan menus offered through proposed exchanges.[1] If consumers do not have the information or abilities to adequately choose an insurance plan, there can be an immediate efficiency loss from consumers not maximizing their individual well-being and a long term efficiency loss from not transmitting the appropriate price signals to the competitive marketplace.

In this work we empirically investigate how one source of choice inadequacy, switching costs, interacts with adverse selection in the context of an employer-sponsored insurance setting typical of the U.S. health care system.[2] Each respective literature has studied these phenomena in isolation from one another: to our knowledge this is the first work that studies them in tandem. In health insurance markets, this interaction matters because choice adequacy impacts plan enrollment, which in turn determines average costs and subsequent premiums. Thus, if there are substantial barriers to decision-making this can have a large impact on the extent of adverse selection and, consequently, consumer welfare. Policies designed to improve consumer choice will have a theoretically 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.

We study individual-level health plan choice and health claims data for the employees of a large firm and their dependents. The data contain a unique natural experiment that we leverage to identify switching costs separately from persistent consumer preference heterogeneity. The firm

[1] Advocates of managed competition in health insurance generally cite policies to help consumers make better decisions as a key component of market regulation (see e.g. Enthoven et al. (2001)).

[2] In 2009, 55.8% of all individuals in the United States (169 million people) received insurance through their employer or the employer of a family member (DeNavas-Walk et al. (2010)). The amount of money at stake in this settting is large: in 2010 the average total premium (employer plus employee contribution) for an employer provided insurance plan was $5,049 for single coverage and $13,770 for family coverage (Kaiser Family Foundation (2010a)).

implemented a major change to their employee insurance program in the middle of the six years of data we observe. The firm substantially changed their menu of five health plan offerings, forced employees out of the health plans they had been enrolled in, and required them to actively choose a plan from the new menu with no stated default option. In subsequent years, the insurance plan options remained the same but consumers had their previously chosen plan as a default option, implying they would continue to be enrolled in that plan if they took no action. This was despite the fact that employee premiums changed markedly over time such that many would have benefited from switching their plan. When combined with other features of the data, our ability to observe the same consumers in clearly active and clearly passive choice environments over time allows us to cleanly identify switching costs.

We present descriptive tests based on the data alone that suggest the presence of large switching costs. Our first test for switching costs studies the behavior of new employees at the firm. As plan prices and the choice environment change over time, incoming cohorts of new employees make active choices that reflect the updated setting while prior cohorts of new employees make markedly different choices that reflect the past choice setup, though they are similar on all other dimensions. A second test studies specific cases that arise in our environment where certain groups of consumers have one of their health plan options become completely dominated by another due to price changes over time. The majority of consumers who face this scenario continue to choose a plan once it becomes dominated, despite the fact that all of them should switch in a frictionless market. Additionally, we present a simple test for adverse selection that clearly reveals that higher health risk employees choose more comprehensive coverage.

While these tests show that switching costs and adverse selection are important in our environment, 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 ex ante health risk. In the model, consumers make choices that maximize their expected utilities over all plan options conditional on their risk tastes and health risk distributions.[3] In the forced active choice periods consumers have zero switching costs, while in periods that they have an incumbent (default) plan option switching costs reduce the utility of alternative options relative to the status quo option. We allow for heterogeneity in both switching costs and risk preferences so that we have the richest possible understanding of how consumers select plans. Though not our primary focus, our risk preference estimates are interesting in their own right as we are aware of only a few previous papers that quantify heterogeneous risk preferences in a non-experimental setting (see e.g. Cohen and Einav (2007) or Chiappori et al. (2008)).

To model health risk perceived by employees at the time of plan choice, we develop a novel out-of-pocket expense model that builds on the prior work of Carlin and Town (2009). The model incorporates past diagnostic and cost information into individual-level and plan-specific expense projections using both (i) sophisticated predictive software developed at Johns Hopkins Medical School and (ii) a detailed model of how different types of medical claims translate to out-of-pocket

[3]In a recent survey of the empirical insurance literature, Einav et al. (2010a) refer to this kind of model that directly estimates expected utility function parameters as a 'realized' utility model.

expenditures in each plan. The cost model outputs a family-plan-time specific distribution of predicted out-of-pocket expenditure distribution that we incorporate into expected utility model under the assumption that consumer beliefs about future health expenditures conform to our cost model estimates.[4] This cost prediction framework directly advances our primary goals of (i) precisely quantifying switching costs and (ii) understanding how plan average costs, and adverse selection, change as enrollment patterns do.

Our choice model estimates reveal that switching costs are high with some substantial heterogeneity, modeled as a function of observable family characteristics. In our primary specification, an average family has switching costs of \$2,032 while the population standard deviation is \$446. An employee covering at least one dependent has, on average, \$751 higher switching costs than a single employee while an employee that enrolls in a flexible spending account (FSA) is estimated to have \$551 lower switching costs than one who does not. These estimates are economically significant relative to the average level of total employee spending of approximately \$4,500. Our risk preference estimates reveal that consumers have a meaningful degree of risk aversion, suggesting that there are, on average, substantial benefits from incremental insurance. We present a variety of robustness analyses to demonstrate that our parameter estimates are quite stable with respect to some of the underlying assumptions in our primary specification.

We apply these estimates to study a counterfactual policy intervention where consumer switching costs are reduced from our baseline estimates. Increased and targeted information provision to consumers is one oft-discussed policy that has the potential to reduce switching costs, though our counterfactual analysis applies to any proposed policies that have the potential to do so.[5] We allow for a range of policy interventions spanning ones that do not change switching costs to those that completely eliminate them. In order to assess the impact of reduced switching costs, it is necessary to model the supply-side of the insurance market. To this end, we construct an insurance pricing model that closely follows the way employee premiums were determined in the firm we study. In our framework, plan premiums equal the average costs of enrollees from the prior period plus an administrative fee, conditional on the number of dependents covered. The firm provides employees with a flat subsidy towards these premiums, implying that consumers pay the full marginal cost of more comprehensive insurance. This pricing environment is very similar to that studied in prior work on insurance markets by, e.g., Cutler and Reber (1998) and Einav et al. (2010b). It also closely resembles the competitive environment of the insurance exchanges recently proposed in the ACA, though there are some specific differences we highlight.

[4]The cost model assumes that there is no (i) consumer-level private information or (ii) moral hazard. While we believe these are potentially important phenomena in insurance markets, in our setting these effects are likely quite small compared to switching costs and selection on the detailed observable information we use. We present a robustness analysis that incorporates estimates from the moral hazard literature to illustrate this point. An earlier version of this paper presented empirical evidence that the combined effect of these two phenomena is not large in our setting (Handel (2010)).

[5]Distinguishing between potential underlying sources of switching costs is important to determine which specific information provision policy will reduce switching costs. In this analysis, we take for granted that there are policies that can reduce switching costs and don't focus on specific policies. Our welfare analysis incorporates the range of potential underlying sources, which we discuss in greater detail in section 7.

In the naive case where plan prices do not change as a result of the different enrollment patterns caused by the intervention, a three-quarter reduction in switching costs substantially improves consumer choices over time. This policy leads to a \$105 (5.2%) mean per person per year welfare increase, measured as the certainty equivalent change as a percentage of employee premiums paid. In the primary policy analysis, where insurance prices endogenously respond to different enrollment and cost patterns, the results are quite different. The same policy that reduces switching costs by three-quarters still improves consumer choices conditional on prices, but now also exacerbates adverse selection, leading to a 7.7% *reduction* in welfare. In this more fluid marketplace, consumers who are healthy and value comprehensive insurance can no longer reasonably purchase it because of the high relative premiums caused by acute sorting. This intervention essentially doubles the existing 8.2% welfare loss from adverse selection in our observed environment, a figure that much of the literature focuses on. We also find, more generally, that welfare is decreasing as the intervention to reduce switching costs becomes more effective. There are substantial distributional consequences resulting from the reduction in switching costs, in addition to the overall efficiency loss. Our welfare analysis accounts for the different potential underlying sources of switching costs by considering a spectrum of cases ranging from the one where switching costs are a true social cost that is fully incorporated into the welfare calculation to the case where switching costs only matter for the choices they engender and do not enter the welfare calculation themselves. The negative welfare impact across the set of policies we consider continues to hold for nearly all welfare treatments of switching costs.

This paper contributes to several distinct literatures. The clean identification of switching costs we obtain from the plan re-design and forced active re-enrollment resolves a primary issue in the empirical literature on switching costs. Farrell and Klemperer (2007) survey this literature and discuss how the inability of researchers to observe active or initial choices within a micro-level panel data set confounds their ability to separately identify switching costs from persistent unobserved preference heterogeneity. The authors note that prior papers that attempt to quantify switching costs generally rely on less direct methods with stronger identifying assumptions. Shum (2004) is a notable example that studies switching costs with panel data in the context of the breakfast cereal markets. Crawford et al. (2011) and Goettler and Clay (2010) are other recent studies that study switching costs with panel data in the fixed-line telephone and grocery delivery markets respectively. There are also extensive empirical marketing literatures on brand loyalty and state dependence, phenomena that are similar in spirit to switching costs (see e.g. Dube et al. (2008) or Dube et al. (2010)).

There is a related literature that studies the effects of inertia in various contexts, such as health insurance (e.g. Strombom et al. (2002) and Ericson (2010)) and 401(k) plan choice (Madrian and Shea (2001)).[6] Our work differs from these papers on several dimensions including that (i) we

[6]Madrian and Shea (2001) cleanly identifies the effect of inertia stemming using a change in the default options for the 401(k) choices of the employees at a large firm. They find that inertia has a substantial impact on 401(k) enrollment and elections. Their work is similar in spirit to our descriptive analysis of new employee health plan choice, though their analysis also expands on this angle to draw a variety of additional conclusions.

quantify switching costs, risk preferences, and ex ante health risk with a choice model and (ii) we study switching costs together with adverse selection in a setting with endogenous pricing. Abaluck and Gruber (2011) use a micro-level data set with plan choices and medical claims in Medicare Part D together with an empirical choice model to show that elderly consumers exhibit specific decision biases that reduce welfare when prices are held fixed, as in our naive model.

This paper builds on the prior work that studies the existence and consequences of adverse selection in health insurance markets. Our insurance choice model relates to the approach of Cardon and Hendel (2001), which is also similar to the approaches used in Carlin and Town (2009), Bundorf et al. (2010), and Einav et al. (2011). These papers model selection as a function of expected health risk and study the welfare loss from adverse selection in their observed settings relative to the first-best. Our work is the first in this literature to identify and estimate switching costs and investigate the interaction between switching costs and adverse selection. In fact, to our knowledge, this is the first paper in the literature to quantify the interaction between any source of choice inadequacy and adverse selection. Additionally, our cost model makes a methodological contribution by showing how to incorporate medical diagnostics and detailed health plan characteristics into individual-level out-of-pocket expense distribution predictions. Using different underlying empirical frameworks, Cutler and Reber (1998) and Einav et al. (2010b) also study the welfare consequences of adverse selection in the context of large self-insured employers. For a more in depth discussion of this literature see the recent survey by Einav et al. (2010a).

Our analysis of switching costs and adverse selection also adds to the empirical literature that studies the impact of insurance selection on preference dimensions separate from risk. Cutler et al. (2008), Cutler et al. (2009), Fang et al. (2008), and Einav et al. (2011) study alternative dimensions of selection in health insurance markets (e.g. risk preferences and moral hazard) while Cohen and Einav (2007) and Einav et al. (2010c) use empirical choice models to study such dimensions in auto insurance and annuity markets respectively. Switching costs are a new potential dimension of selection that differs from those previously studied because it leads to specific selection patterns as consumers and the marketplace change over time (as opposed to a purely static context).

The rest of the paper proceeds as follows. Section 2 describes the data with an emphasis on how the health insurance choice environment evolves at the firm over time. Section 3 presents descriptive tests based on the data alone that show the presence of both switching costs and adverse selection. Section 4 presents our empirical choice framework while section 5 presents the structural estimates from this model. Section 6 presents a model of insurance pricing, describes our welfare framework, and investigates the impact of counterfactual policies that reduce switching costs relative to our estimates. Section 7 discusses and section 8 concludes.

2 Data and Environment

We study the health insurance choices and medical utilization of employees (and their dependents) at a large U.S. based firm over the time period from 2004 to 2009. In a year during this period that we denote t_0 (to protect the identity of the firm) the firm changed the menu of health plans it offered to employees and introduced an entirely new set of PPO plan options.[7] At the time of this change, the firm forced all employees to leave their prior plan and actively re-enroll in one of five options from the new menu, with no default option. The firm made a substantial effort to ensure that employees made active choices at t_0 by continuously contacting them via physical mail and email to communicate both information about the new insurance program and that there would be no default option.[8] In the years prior to and following the active choice year t_0, employees were allowed to default into their previously chosen plan option without taking any action, despite the fact that in several cases plan prices changed significantly from one year to the next. This variation in the structure of the default option over time, together with the plan menu change, is a unique feature of the data set that makes it especially well suited to study switching costs because, for each longer-term employee, we observe at least one choice where switching costs could be present and one choice where they are not.

These proprietary panel data include the health insurance options available in each year, employee plan choices, and detailed, claim-level employee (and dependent) medical expenditure and utilization information.[9] We use this detailed medical information together with medical risk prediction software developed at Johns Hopkins Medical School to develop individual-level measures of predicted future medical utilization at each point in time. These measures are generated using past diagnostic, expense, and demographic information and allow us to precisely gauge medical expenditure risk at the time of plan choice in the context of our cost model. [10] Additionally, we observe a rich set of employee demographics including job characteristics, age, gender, income, and job tenure, along with the age, gender, and type of each dependent. Together with data on other relevant choices (e.g. flexible spending account (FSA) contributions, dental insurance) we use these characteristics to study switching cost and risk preference heterogeneity.

Sample Composition and Demographics. The firm we study employs approximately 9,000 people per year. The first column of Table 1 describes the demographic profile of the 11,253 employees who work at the firm for some stretch within 2004-2009. These employees cover 9,710

[7] This change had the two stated goals of (i) encouraging employees to choose new, higher out-of-pocket spending plans in order to help control total medical spending and (ii) providing employees with a broader choice of different health insurance options (e.g. a consumer driven health plan with a linked health savings account (HSA)).

[8] Eventually, though they were not told this ahead of time, the approximately 0.6% employees that did not actively elect a given plan were all enrolled in one option so that they would not lose out on this valuable benefit.

[9] We observe this detailed medical data for all employees and dependents enrolled in one of several PPO options, which is the set of available plans our analysis focuses on. For a further discussion see the sample composition section. These data include detailed claim-level diagnostic information, such as ICD-9 codes and NDC drug codes, as well as provider information and a detailed payment breakdown (e.g. deductible paid, coinsurance paid, plan paid).

[10] This program, known as the Johns Hopkins ACG (Adjusted Clinical Groups) Case-Mix System, is one of the most widely used and respected risk adjustment and predictive modeling packages in the health care sector. It was specifically designed to use diagnostic claims data, such as the individual-level ICD-9 codes we observe, to predict future medical expenditures in a sophisticated manner. In addition, the program takes into account the NDC pharmaceutical drug utilization codes we observe as well as individual age and gender. For more information on this program visit *http://www.acg.jhsph.edu/html/AboutACGs_whois.htm*. Other recent work by Carlin and Town (2009), Bundorf et al. (2010), and Einav et al. (2011) uses similar software.

dependents, implying a total of 20,963 covered lives. 46.7% of the employees are male and the mean employee age is 40.1 (median of 37). We observe income grouped into five tiers, the first four of which are approximately $40,000 increments, increasing from 0, with the fifth for employees that earn more than $176,000. Almost 40% of employees have income in tier 2, between $41,000 and $72,000, with 34% less than $41,000 and the remaining 26% in the three income tiers greater than $72,000. 58% of employees cover only themselves with health insurance, with the other 42% covering a spouse and/or dependent(s). 23% of the employees are managers, 48% are white-collar employees who are not managers, and the remaining 29% are blue-collar employees. 13% of the employees are categorized as 'quantitatively sophisticated' managers. Finally, the table presents information on the mean and median characteristics of the zip codes the employees live in.

We construct our final sample to leverage the features of the data that allow us to identify switching costs. Moving from the full data, we restrict the final sample to employees (and dependents) who (i) are enrolled in a health plan for all years from t_{-1} to t_1 and (ii) are enrolled in a PPO option in each of those years (excluding the employees that enroll in either of two HMO options).[11] The second column in Table 1 describes the sample of employees who ever enroll in a PPO option at the firm (N = 5,667), while the third column describes the final sample (N = 2,023). Comparing column 2 to column 1, it is evident that the restriction to PPO options engenders minimal selection based on the rich set of demographics we observe. Comparing both of these columns to column three reveals that the additional restriction that employees be enrolled for three consecutive years does lead to some sample selection, though not a substantial amount. Employees in the final sample are slightly older, slightly richer, and more likely to cover additional family members than the overall $PP0$ population. Note that the primary impact of the restriction that employees be enrolled the entire three year period is to exclude employees who enter or exit the firm during that period, rather than exclude employees who switch to an HMO option or waive coverage.

There are costs and benefits of these two restrictions. The restriction to PPO plans is advantageous because we observe detailed medical claims data only for enrollees in these plans and these plans are only differentiated by financial characteristics, implying we don't have to consider heterogeneity in preferences over provider network when modeling choice between them.[12] A potential cost is that this restriction may bias the choice model by restricting the set of options. In the upcoming descriptive analysis of plan choices we show clear evidence that the nest of PPO options and nest of HMO options are quite horizontally differentiated from one another, implying a limited within-sample bias from excluding HMO choices.

The restriction that employees enroll in a plan in every year from t_{-1} to t_1 has the benefit that, for each individual in the final sample, we observe a past year of medical data for each choice spanning t_0 (the active choice year) to t_2. This allows us to model health risk at the time of each choice from an ex ante perspective, permitting a more precise characterization of out-of-pocket expense

[11] We denote all years in reference to t_0, such that, e.g., year t_{-1} occurred just before t_0 and year t_1 just after.

[12] This implies that, ultimately, the switching costs that we study exclude potential switching costs that come from having to change your medical provide when you move to a new health plan. Thus, our estimates can be interpreted as a lower bound on switching costs in a more general setting where this factor is potentially important.

risk and the choice model parameters. This restriction has two costs: (i) it reduces the sample size and (ii) it excludes new employees from t_0 to t_2, who, as the upcoming preliminary analysis section reveals, can provide an additional source of identification for switching costs.[13] Ultimately, since the identification within the final sample for switching costs is quite strong because of the plan menu change and linked active decision, we feel that having a more precise model is worth the costs of this restriction.[14]

Health Insurance Choices. From 2004 to t_{-1} the firm offered five total health plan options composed of four HMO plans (smaller provider network, more integration with providers) and one PPO plan (broader network, less integration). Each of these five plans had a different network of providers, different contracts with providers, and different premiums and cost-sharing formulas for enrollees. From t_0 on, the new plan menu contained two of the four incumbent HMO plans and three new PPO plans.[15] This plan structure remained intact through the end of the data in 2009. After the menu change, the $HMOs$ still had different provider networks and cost sharing rules both relative to each other and to the set of new $PPOs$. However, as mentioned, the three new PPO plans introduced at t_0 had exactly the same network of providers, the same contractual treatment of providers, and cover the same medical services. The PPO plans are only differentiated from one another (and from the previously offered PPO) by premiums and cost sharing characteristics that determine the mapping from total medical expenditures (employee paid plus employer paid) to employee out-of-pocket expenditures (e.g. deductible, coinsurance, and out-of-pocket maximums). Throughout the period, all PPO options that the firm offers are self-insured plans where the firm fills the primary role of the insurer and is at risk for incurred claims. We denote the HMO plans available throughout the entire period as HMO_1 and HMO_2, and those offered only before the menu change as HMO_3 and HMO_4. We denote the PPO option from before the menu change as PPO_{-1}, while we we denote each of the PPO options after the menu change by their respective individual-level deductibles: PPO_{250}, PPO_{500} and PPO_{1200}.[16] PPO_{1200} is paired with a health savings account (HSA) option that allows consumers to deposit tax-free dollars to be used later to pay medical expenditures.[17]

[13]Note also that this restriction excludes employees who exit before t_1, but that these employees would only have one choice in the active choice year in the model, making them less useful for identification of switching costs.

[14]If we included new employees after t_{-1} 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, similarly to what is done in the rest of the literature when detailed claims information is not available.

[15]Any employee who chose a PPO plan at t_0, by definition, had to be actively choosing an entirely new plan. An employee was enrolled in HMO_1 or HMO_2 prior to t_0 was still forced out of that plan and prompted to make an active choice from the new menu, though those plans remained available. Since we only study PPO plans after t_0, the incumbent aspect of the HMO plans does not impact our analysis.

[16]The deductibles can be linked directly to how comprehensive (proportion of health expenses paid for) each plan is overall. Thus, PPO_{250} is the most comprehensive (and highest premium) option, PPO_{500} is in the middle, and PPO_{1200} is the least comprehensive (lowest premium).

[17]This may lead to some degree of horizontal differentiation for this plan relative to the other two, which we account for in the choice model. This kind of plan is known as a 'high-deductible health plan' or 'consumer driven health plan'. Employees who signed up for this plan for the first time were given up to a $1,200 HSA match from the firm, which we account for in our analysis.

Table 2 presents the detailed characteristics of the PPO plans offered after the menu change at t_0. After the deductibles are paid, PPO_{250} has a coinsurance rate of 10% while the other two plans have rates of 20%, implying they have double the marginal price of post-deductible claims. Out-of-pocket maximums indicate the maximum amount of medical expenditures that an enrollee can pay post-premium in a given plan. These are larger the less comprehensive the plan is and vary with income tier. Finally, both PPO_{250} and PPO_{500} have the same flat-fee copayment structures for pharmaceuticals and physician office visits, while in PPO_{1200} these apply to the deductible and coinsurance.[18] While we model these characteristics at a high-level of detail, our cost model necessarily makes some simplfying assumptions that we discuss and validte in Appendix A.

The top panel in Figure 1 compares plans PPO_{250} and PPO_{500} graphically, to illustrate the relationship between health plan financial characteristic and employee out-of-pocket expenses. The figure studies this relationship at t_0 (premiums differ each year) and describes employee out-of-pocket expenditures as a function of total medical expenses that apply to the deductible, coinsurance and out-of-pocket maximums. The figure applies to low income families but looks similar in structure for other coverage tiers and income levels. The figure completely represents the difference between these two plans, since the plans are identical on the co-payments for pharmaceuticals and office visits excluded from this chart. Throughout our analysis we assume that (i) premiums are in pre-tax dollars and (ii) medical expenses are in post-tax dollars.[19] After the employee premium (vertical intercept), as total expenditures increase each employee pays the plan deductible, then the flat coinsurance rate, and finally has zero marginal cost after reaching the out-of-pocket maximum.[20] As expected, the chart reveals that, from an ex post perspective, healthy employees shuold have chosen PPO_{500} and sick employees PPO_{250}.

The top part of Table 3 shows the pattern of employee choices over time before and after the menu change. In t_{-1}, 39% of employees enroll in PPO_{-1}, 47% enroll in one of the four HMO options, and 14% waive coverage. At t_0, 46% of employees choose one of the three new PPO options, with 25% choosing PPO_{250}. 37% choose either of the two remaining HMO plans while 16% waive coverage. The bottom two parts of table 3 study plan choice transitions over time, and

[18]These characteristics are for in-network purchases. We note that the plans also have slightly different out-of-network characteristics, which we do not present or model. The plans have reasonably similar characteristics (including out-of-pocket maximums) for out-of-network claims while only 2% of realized claims expenditures apply to this domain.

[19]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 believe this is the correct empirical assumption. Moreover, the connected paper Handel et al. (2011) reveals that consumers should optimally make FSA contributions much less than their expected expenditures. 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. These adjustments do take into account marital status and family size but do not take into account spouse's income, for which we have no available data. As a result, we may understate the marginal tax rate for employees with high-earning spouses.

[20]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.

present clear evidence that the nest of PPO options and nest of HMO options are quite horizontally differentiated from one another. An individual who switches plans from a PPO option is much more likely to choose another PPO option than to choose an HMO option. The middle panel shows that, of the 2,757 employees enrolled in PPO_{-1} in year t_{-1} who also enroll in any plan at t_0, only 85 (3%) choose an HMO option at t_0. In reverse, despite the expansion of PPO options and reduction of HMO options, only 15% of employees who chose an HMO option in t_{-1}, and choose any plan at t_0, switch to a PPO option. This is strong evidence that restricting the set of choices to PPO options should not lead to biased parameters within that population.[21]

Each plan offered by the firm has a distinct 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 amount the employee actually pays after receiving a subsidy from the firm.[22] Total premiums are conditioned on being in one of four coverage tiers.[23] The firm conditions PPO subsidies on an employee's income tier, presumably because of equity considerations.[24] Figure 2 illustrates employee premium contributions in years t_0 and t_1 for the single and family (spouse plus children) coverage tiers. There is a noticeable decrease in premiums for PPO_{500} from t_0 to t_1 coupled with an increase in the premium for PPO_{250}. For example, for a family in the top income tier, the price of PPO_{500} decreased by $\$1,560$ from t_0 to t_1 while the price of PPO_{250} increased by \$420.[25] There are also substantial relative premium changes for the other three coverage tiers. As a result of these large relative employee premium changes, 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.

3 Preliminary Analysis

We start the analysis by presenting some descriptive evidence of switching costs and adverse selection. We investigate two different model-free tests that suggest switching costs are an important factor in determining choices over time. In addition, we present a test for adverse selection based on the data alone. While this section presents strong evidence on the existence and potential impact of these two phenomena, it also highlights that a more in depth modeling exercise is essential to precisely quantify the magnitudes of these effects and evaluate the impact of a counterfactual reduction in switching costs. Each analysis uses a sample that differs from our primary sample because of the specific source of identification involved.

[21] Einav et al. (2011) find similar horizontal differentiation in a similar context.

[22] 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. While in theory these self-insured total premiums could be set in a variety of different ways to reflect different distributional aims, in our setting total premiums are set in a specific way that is similar both to other large employers and to what we would expect to arise in certain quasi-private markets such as health insurance exchanges. This is discussed at length in our pricing analysis in section 6.

[23] These are (i) single (ii) employee + spouse (iii) employee + child(ren) and (iv) employee + spouse + child(ren).

[24] The firm gives employees a lump sum subsidy that applies to each potential PPO option. The firm sets a target of subsidizing 76% of total premium payments for employees.

[25] This movement is due to total premium adjustment based on t_0 average costs for each plan and dependent coverage tier and reflects incremental adverse selection against PPO_{250}.

New Employees. Our first test for switching costs studies the behavior of new employees at the firm over time. New employees are an interesting group to investigate because they have zero switching costs when they choose a new health plan at the time of their arrival. This is because (i) they have no health plan default option at the time of arrival and (ii) they were not previously enrolled in any health plan within the firm.[26] Thus, in our setting, employees who are new for year t_0 have zero switching costs in that period and positive switching costs when choosing a plan for year t_1 (this is the same choice structure over time as existing employees, given the menu change). Moving forward, employees who are new in year t_1 have no switching costs at t_1 and positive switching costs thereafter. Given the large price changes for t_1 described in the prior section, if the profile of new employees is similar in each year then large switching costs should imply that the t_1 choices of new enrollees at t_0 are different than the t_1 choices of new enrollees at t_1. In that case, the t_1 choices of t_0 new enrollees should reflect the choice environments at both t_0 and t_1, while the t_1 choices of t_1 new enrollees should depend on just the t_1 environment.

Table 4 compares the choices over time of the cohorts of new enrollees from years t_{-1}, t_0, and t_1, with each group composed of slightly more than 1,000 employees. Without switching costs, we would expect the choices in these three cohorts to be the same at t_1, since the table reveals that they are virtually identical on all other demographic dimensions, including age, gender, income, FSA enrollment, and health expenditures. Instead, while it is evident that the t_0 and t_{-1} cohorts make very similar choices with the default option at t_1, the new enrollees making active choices in that year have a very different choice profile that reflects the price changes for t_1. For example, 21% of t_0 new enrollees choose PPO_{250} at t_0 while 23% choose PPO_{500}. At t_1, 20% of this cohort choose PPO_{250} and 26% choose PPO_{500} only a small change in market share for each plan in the direction expected given the price changes. The decision profile over time for new enrolles at t_{-1} is similar. However, new enrollees at t_1 choose PPO_{250} only 11% of the time, while choosing PPO_{500} 43% of the time. This implies that t_0 and t_{-1} new 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.

Dominated Plan Choice. Our second test for switching costs leverages a unique situation caused by the combination of plan characteristics and plan price changes in our setting. As a result of the large price changes for year t_1, PPO_{250} became *strictly dominated* for certain combinations of family size and income (recall, this is how employee premium contributions are determined). Strict dominance implies that for any possible level and type of total medical expenditures, PPO_{500} leads to lower employee expenditures (premium plus out-of-pocket) than PPO_{250}. The bottom panel in Figure 1 reproduces, for year t_1, the t_0 analysis of PPO_{250} and PPO_{500} health plan characteristics discussed earlier. The figure studies the relationship between total medical expenditures (employee

[26] Moreover, because each of the PPO options we study has the exact same network of providers, there is no built-in advantage for specific plans because of prior coverage. Further, since the PPO options are self-insured, these specific plans are not offered in the same names and formats at other firms or in the private market.

and insurer) and employee expenditures (premium plus out-of-pocket) for low income families (employee plus spouse plus dependent(s)). For this group, the large relative premium change between these two plans for t_1 shifts the relative baseline employee expenditures so much that a low income family should always enroll in PPO_{500} at t_1 if making an active choice, regardless of beliefs about future medical expenditures. In fact, the figure illustrates that a low income family that enrolls in PPO_{250} at t_1 loses at least \$1,000 for sure relative to PPO_{500}. Recall that this chart represents all dimensions of differentiation between these two plans. At t_0, with the active re-enrollment, there were no dominated plans for any employee. PPO_{250} is dominated at t_1 for four of the other nineteen potential coverage and income tier combinations. It is important to note that the existence of dominated plans for these select groups was unknown to the firm at t_1. Total premiums and employee subsidies were determined completely separately from the decision on health plan characteristics made at t_0, implying that the firm did not analyze these features in combination with each other at t_1 and t_2 as we do here.

Table 5 describes the behavior of the subset of employees who enrolled in PPO_{250} at t_0 and had that plan become dominated for them in t_1 and t_2. Of the 1,897 employees who enroll in PPO_{250} at t_0 and remain with the firm at t_1, 559 (29%) had that plan become dominated for them in t_1 (for t_2 504 of these remain at the firm). Of these 559 employees, only 61 (11%) switch plans to an undominated plan at t_1 indicating substantial persistence in plan choice that must be the result of switching costs (under a broad definition) because unobserved preference heterogeneity cannot rationalize choosing PPO_{250} at t_1. Thus, for these employee groups, in a rational frictionless environment we would expect 100% of the individuals enrolled in PPO_{250} at t_0 to switch to PPO_{500} at t_1. Of the 61 employees that did switch at t_1, the majority (44 (72%)) switch to PPO_{500} as expected given an active choice and the large relative price drop of that plan. This pattern remains similar even at t_2 after employees have had more time to communicate with one another: only 126 (25%) of the 504 employees switch by t_2, with 103 (82%) switching to PPO_{500}. The table reveals that the average minimum money lost by employees in these groups from staying in PPO_{250} is \$374 at t_1 and \$396 at t_2.

Table 5 also reveals that employees who switch plans over time are more likely to make other active decisions. The top part of the table describes linked FSA and dental plan decisions for those with dominated plans, while the bottom panel describes these choices for people who switch from any PPO option in the entire population. Conditional on switching from a dominated option at t_1, 14.1% of employees also switch dental plan at t_1, compared to 4.3% for those who don't switch. For the entire population and universe of PPO plans (3,170 employees present over multiple years) the analagous numbers switching dental plan are 14.5% and 3.8%. The table also reveals that employees who switch plans at t_1 are more likely to enroll in an FSA at t_1. This is a relevant choice to study because FSA enrollment is an active choice in each year: employees who do not actively elect to sign up and list a contribution level are not enrolled. For the entire population in PPO plans, 25% of those who do not switch sign up for an FSA at t_1 while 39% of those who do switch sign up (the pattern is similar for those with dominated plan options). This correlation could indicate either that (i) employees who enroll in and FSA are more generally active consumers or that (ii) when they make the active choice to switch health plan this causes them to also actively enroll in an FSA. The analysis also reveals that those who switch are, on average, younger, slightly lower income, and more likely to be male.

Adverse Selection. Before we present the main econometric framework, we provide evidence that some adverse selection is present in the data we observe. Table 6 studies the choice and cost behavior of our primary sample described in the previous section. The top panel shows the level of t_{-1} claims for individuals enrolled in each of the PPO options from t_{-1} to t_1. We study t_{-1} claims for plans chosen across all three years t_{-1} to t_1 specifically to avoid the potential alternative explanation of moral hazard: in year t_{-1} all families in this sample were enrolled in PPO_{-1} implying that t_{-1} claims are an 'apples to apples' measure of health expense risk that is not confounded by moral hazard. The table reveals that there is selection on medical expenses against the most comprehensive plan, PPO_{250}. Employees who chose PPO_{250} had almost double the median and mean of t_{-1} total medical claims relative to enrollees in the other two PPO options, in both t_0 and t_1. Despite the large price change from t_0 to t_1, the pattern of selection barely changes over these years. The high level of selection at t_0 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 over time, even though prices changed significantly. This suggests that high switching costs likely reduce adverse selection in t_1 relative to what it would have been if everyone had made active plan choices in that year. This motivates our counterfactual exercise investigating the impact of an information provision policy that reduces switching costs in the context of a setting with adverse selection.

4 Empirical Framework

The analysis in the previous section provides evidence of both large switching costs and adverse selection without imposing specific choice and cost models. This section presents a model of consumer choice with three primary components (i) switching costs (ii) family-level risk preferences and (iii) individual-level *ex ante* cost projections. We describe the empirical implementation of this model, which links the choice and medical cost data we observe to these underlying economic choice fundamentals. Relative to the earlier analysis, this framework makes it possible to (i) quantify switching costs and (ii) determine the impact of potential counterfactual information provision policies that reduce switching costs. These additional conclusions should be viewed in the context of the structural assumptions included in the model. We present the supply-side insurance pricing model later in section 6, together with the analysis of the interaction between switching costs and adverse selection.

Choice Model. We describe the model in two components. First we describe the choice framework *conditional* on predicted family-level ex ante medical cost risk. Afterwards, we describe the detailed cost model that generates these cost distributions.

The choice model quantifies switching costs and risk preferences conditional on the family-plan-time specific distributions of out-of-pocket health expenditures output by the cost model. Denote these expense distributions $F_{kjt}(\cdot)$, where k is a family unit, j is one of the three PPO insurance plans available after the t_0 menu change, and t is one of three years from t_0 to t_2. We assume that families' beliefs about their out-of-pocket expenditures conform to $F_{kjt}(\cdot)$. 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 et al. (2010a) call a 'realized' empirical utility model and assume that U_{kjt} has the following von-Neuman Morgenstern (v-NM) expected utility formulation:

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

Here, $u_k(\cdot)$ is the v-NM utility index and OOP is a realization of medical expenses from $F_{kjt}(\cdot)$. W_k denotes family-specific wealth. P_{kjt} is the family-time specific premium contribution for plan j, which as described earlier depends both on how many dependents are covered and on employee income.[27] $\mathbf{1}_{kj,t-1}$ is an indicator of whether the family was enrolled in plan j in the previous time period.

We assume that families have constant absolute risk aversion (CARA) preferences implying that for a given ex post consumption level x:

$$u_k(x) = -\frac{1}{\gamma_k(X_k^A)}e^{-\gamma(X_k^A)x}$$

Here, γ_k is a family-specific risk preference parameter that is known to the family but unobserved by the econometrician. We model this as a function of employee demographics X_k^A. 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.[28]

In our primary empirical specification a family's overall level of consumption x conditional on a draw OOP from $F_{kjt}(\cdot)$ depends on multiple factors:

$$x = W_k - P_{kjt} - OOP + \eta(X_{kt}^B, Y_k)\mathbf{1}_{kj,t-1} + \delta_k(Y_k)\mathbf{1}_{1200} + \alpha H_k\mathbf{1}_{250} + \epsilon_{kjt}(Y_k)$$

η is a switching cost that depends on the observable linked choice and demographic variables X_{kt}^B and Y_k. We describe these in more detail in the estimation section. δ_k is an unobserved family-specific plan intercept for PPO_{1200} (1_{1200} is an indicator for $j = PPO_{1200}$). On average, we expect δ_k to differ from zero because the health savings account (HSA) option offered exclusively through PPO_{1200} horizontally differentiates this plan from the other two PPO options.[29] α measures

[27]We treat premiums as pre-tax payments and out-of-pocket expenditures as after-tax. See our data description in section 2 for a further discussion.

[28]This implies that wealth W_k does not impact relative plan utilities. As a result, it drops out in estimation. The specification for wealth would matter under an alternative framework such as constant relative risk aversion (CRRA) preferences.

[29]Prior research shows that HSAs can cause significant hassle costs or, alternatively, provide an extra benefit in the

the intrinsic preference of a high-cost family for PPO_{250}, where high-cost, represented by the binary variable H_k, is defined as greater than \$27,000 (approximately 90th percentile of total cost distribution) in any of the three years in the sample.[30] Finally, ϵ_{kjt} represents a family-plan-time specific idiosyncratic preference shock. Since the plans we study are only differentiated by financial characteristics (apart from the HSA feature) we also follow Einav et al. (2011) and study a robustness check later with no idiosyncratic preference shock.

There are several assumptions in the choice model that warrant additional discussion. The model assumes that families know the distribution of their future health expenditure risk and that this risk conforms to the output of the cost model described in the next section. This assumption could be incorrect for two potential reasons. First, families may have private information about their health status that is not captured in the detailed prior claims data. Second, families may have *less* information about their projected future health expenditures. As described shortly, the cost model utilizes a full profile of past claims data in conjunction with sophisticated software that maps past claims to future expected expenses. Further, the model contains the assumption that consumers have full knowledge of each health plan's characteristics and incorporate that knowledge into their decision process. Each of these potential deviations implies a potential bias in the F_{kjt} distributions. We present a robustness analysis with our results to show that reasonable sized deviations from our estimates of F_{kjt} do not substantially affect the estimated switching costs or other choice model parameters.

Additionally, switching costs are modeled in a specific way, as an incremental cost paid conditional on actually switching plans (which could arise from a variety of micro-foundations). 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 follows the approach used in the theoretical literature (see e.g. Farrell and Klemperer (2007)) as well as in the prior empirical work on switching costs (see e.g. Shum (2004) or Dube et al. (2008)). After we present our results, we discuss a variety of potential different sources of the switching costs we find (e.g. transaction costs, learning, rational inattention, and inertia) along with several corresponding alternative specifications. While these alternative specifications would capture the evident persistence in plan choice with different underlying mechanisms, we argue in the discussion that their implications for how switching costs interact with adverse selection would not differ substantially from our primary model.[31]

Further, our model assumes that consumers are myopic and do not make dynamic decisions

form of an additional retirement account (see e.g. Reed et al. (2009) or McManus et al. (2006)). Further, consumer uncertainty about how HSAs function could deter choice of the high-deductible health plans linked directly to HSAs. Though we have data on exact family HSA contributions and linked employer contribution matches for first-time enrollees, we choose to incorporate preferences for these benefits into the coefficient δ_k in lieu of a more detailed model.

[30] This latter variable is included to proxy for the fact that almost all 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.

[31] Also, when we study the welfare consequences of selection resulting from reduced switching costs, we present a range of welfare calculations that span how switching costs themselves are treated, which links back to these different potential sources.

whereby current choices would take into account switching costs in future periods. There are several arguments to support this approach. First, price changes are not signaled in advance and change as a function of factors that would be difficult for consumers to model.[32] Second, it is unlikely that most consumers can forecast substantial changes to their health status more than one year in advance. Third, in this empirical setting consumers make initial choices that make little sense in the context of a fully dynamic approach. They choose (and stay with) plans at t_0 that provide poor long run value given the time path of prices and health expectations. While it is clear that dynamics are an important component of choice with switching costs in some settings (e.g. cable provider or cell phone plan choice), in this setting we believe that incorporating empirically relevant dynamic aspects would not markedly change the results.

Cost Model. The choice framework presented in the previous section takes the distribution of future out-of-pocket expenditures for a given family, $F_{kjt}(\cdot)$, as given. This section describes the empirical model we use to estimate $F_{kjt}(\cdot)$, at a high level. Appendix A delves into this material in more detail and presents a formal description of the model, its estimation algorithm, and its output results.

Our approach models health risk and out-of-pocket expenditures at the individual level, and aggregates the latter measure to the family level since this is the relevant metric for plan choice. For each individual and choice period, we model the distribution of future health risk at the time of plan choice using past diagnostic, demographic, and cost information. This ex ante approach to the cost model fits naturally with the insurance choice model where families make plan choices under uncertainty. In the majority of prior work investigating individual-level 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 at the individual level, from past or futures years (Carlin and Town (2009), Einav et al. (2011), and Abaluck and Gruber (2011) are notable exceptions). While these approaches are useful approximations when detailed medical data are not available, our model is able to more precisely characterize a given family's information set at the time of plan choice and can be linked directly to the choice problem.

Our model makes several advances relative to the recent literature that uses micro-level claims data to quantify individual health risk. First, we use past diagnostic data, in combination with the Johns Hopkins ACG software package, to predict total future medical expenditures in a sophisticated manner. This method incorporates, for example, the duration of conditions and the implications of multiple simultaneous diagnoses.[33] Second, we develop a new parsimonious method to non-parametrically link expected future health risk to expected future expenditures by com-

[32] For example, for consumers to understand the evolution of prices they would have to (i) have knowledge of the pricing model (ii) have knowledge about who will choose which plans and (iii) have knowledge about other employees' health.

[33] The program was specifically designed to use diagnostic claims data to predict future medical expenditures in a sophisticated manner. For example, in our model, a 35 year old male who spent \$10,000 on a chronic condition like diabetes in the past year would have higher predicted future health expenses than a 35 year old male who spent \$10,000 in the past year to fix time-limited acute condition, such as a broken arm.

bining the predictive ACG health risk output with observed cost data. Finally, we reconstruct the mapping from total medical expenditures to health plan out-of-pocket costs using a detailed breakdown of medical claims and plan characteristics. To this end, we predictively model health risk across mutually exclusive and exhaustive categories of medical expenditures, including cross-category correlated risks, and link these categories to plan-specific out-of-pocket cost determinants. As a result of these advances, our cost model is unique in the literature for its ability to precisely measure out-of-pocket expense risk. This, in turn, has clear positive implications for the viability and precision of our choice model parameters and subsequent counterfactual analysis. Appendix A studies the intricacies of the cost model in more detail.

The cost model assumes that there is no moral hazard (i.e. total family expenditures do not vary with j) and that there is no family-specific private information. While both of these phenomena have the potential to be important in health care markets, and are studied extensively in other research, we believe that these assumptions do not materially impact our results. One primary reason is that both effects are likely to be quite small relative to the magnitude of switching costs we estimate, which is generally above a thousand dollars. For private information, we should be less concerned than prior work because our cost model combines detailed individual-level prior medical utilization data with sophisticated medical diagnostic software. 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] For moral hazard, Chandra et al. (2010) presents a recent review of the experimental and quasi-experimental literature, where the price elasticity generally falls in the range -0.1 to -0.4. Recent work by Einav et al. (2011) on data similar to that use here finds an implied elasticity of -0.14. We perform an in depth robustness analysis in the next section that incorporates these elasticity estimates into our cost model estimates to verify that the likely moral hazard impact (i) is small relative to the degree of switching costs we measure and (ii) does not markedly impact our parameter estimates.[35]

Identification. Our primary identification concern is to separately identify switching costs from persistent unobserved preference heterogeneity. Prior studies seeking to quantify switching costs have been unable to cleanly distinguish between these phenomena primarily because, in their respective settings, they (i) do not observe periods where consumers make identifiably 'active' choices in some periods and identifiably 'passive' choices in others and (ii) the products in question are differentiated such that persistent consumer preference heterogeneity is a distinct entity for each product. For a further recent discussion of these issues see, e.g., Dube et al. (2010).

[34]Cardon and Hendel (2001) find no evidence of selection based on private information with more coarse data. Pregnancies, genetic pre-dispositions, and non-coded disease severity are possible examples of private information that could still exist (technically we could 'back date' pregnancies we observe later in the data to control for this source). Carlin and Town (2009), whose cost model is done with detailed medical information, also argue that significant residual selection is unlikely. Importantly, it is also possible that individuals know *less* about their risk profile than we do, since we have a detailed claims record and sophisticated medical software.

[35]Handel (2010), a prior version of this paper, presents descriptive empirical evidence that the combined effects of selection on private information and moral hazard are not large in our setting. This is similar in spirit to the work of Chiappori and Salanie (2000).

Our setting is ideally suited to overcome these difficulties and identify switching costs. First, the plan menu change and forced re-enrollment at year t_0 ensures that we observe each family in our final sample making both an 'active' and a 'passive' choice from the same menu of health plans (same except for year on year price changes). Second, the three PPO plan options we study have the exact same network of medical providers and cover the same medical services, implying that differentiation occurs only through preferences for plan financial characteristics (here, risk preferences). Third, since insurance choice here is effectively a choice between different financial lotteries, our detailed medical data allow us to precisely quantify health risk and the ex ante value consumers should have for health plans, conditional on the assumption that beliefs conform to F_{kjt} and on their risk preferences. Finally, we study three consecutive choices where prices change substantially over time, which allows us to see consumers making choices in environments where many *should* switch absent switching costs. The combined effect of these features implies that (i) we identify consumer preference heterogeneity based on the choices made in the active choice, forced re-enrollment period, while (ii) switching costs are identified by analyzing how choices change over time as the predicted active plan values change. Thus, if switching costs were zero, t_1 choices with the changed plan valuations would reflect the preferences identified at t_0. Positive switching costs imply that t_1 choices will reflect the old choice environment at t_0 as well as the new environment in t_1. We can then quantify switching costs by assessing how much value is being foregone at t_1 through the continued impact of t_0 choices.[36]

The random coefficient δ captures an additional preference for PPO_{1200} resulting from the linked HSA option. δ is identified separately from the family-specific risk aversion parameter, γ, by leveraging the structure of the three available choices. γ is identified separately by the choice between PPO_{250} and PPO_{500}, which are not horizontally differentiated in any way, and δ_k is then identified by examining the choice between the nest of those two plans and PPO_{1200}. These two sources of preference heterogeneity are then identified separately from switching costs as described above.

Estimation. In our primary specification, we assume that the random coefficients γ_k and δ_k are normally distributed with a mean that is linearly related to relevant observable characteristics X_k^A:[37]

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

In the primary specification X_k^A contains employee age and income.[38] We also investigate a robustness check with log-normally distributed γ. We denote the mean and variance of δ_k, the random

[36]Our earlier preliminary analysis reveals that studying the behavior of new employees over time is another avenue to identify switching costs.

[37]For normally-distributed γ, we assume that γ is truncated just above zero.

[38]While age and income do change over time, they vary minimally over the three-year estimation period so we treat them as constant for a given individual. These represent the individual employee whose family is choosing insurance.

intercept for PPO_{1200}, as $\mu_\delta(Y_k)$ and $\sigma^2_\delta(Y_k)$. These quantities are estimated conditional on the binary family status indicator Y_k, with the two categories of (i) single and (ii) family covering dependents.[39]

Switching costs, $\eta(X^B_{kt}, Y_k)$, are related linearly to Y_k and observable linked choices and demographics X^B_{kt}:

$$\eta(X^B_{kt}, Y_k) = \eta_0 + \eta_1 X^B_{kt} + \eta_2 Y_k$$

X^B_{kt} contains potentially time-varying variables that switching costs may depend on, including income and whether or not (i) the family enrolls in a flexible spending account (FSA) (ii) the employee has a quantitative background (iii) the employee is a manager within the firm (iv) a family member has a chronic medical condition (v) the family has a large change in expected expenditures from one year to the next or (vi) the family switches away from PPO_{1200}. Note here that many of the X^B_{kt} and Y_k conditioning variables are binary, implying the linearity assumption is not restrictive for these variables.

Finally, we assume that the probit family-plan-time specific error terms ϵ_{kjt} are distributed i.i.d. for each j with mean $\mu_{\epsilon_j}(Y_k)$ and variance $\sigma_{\epsilon_j}(Y_k)$. Since Y_k is a binary variable we make no additional assumptions on how these means and variances relate to Y_k. We normalize the value of ϵ_{250}, the preference shock for PPO_{250}, to zero for each realization of Y_k, and estimate the preference shock means and variances for the other two plans relative to PPO_{250}. [40] Since the set of PPO plans we study can be compared purely based on financial characteristics (conditional on the already modeled HSA option), we follow Einav et al. (2011) and study a robust specification without ϵ_{kjt}.

We estimate the choice model using a random coefficients probit simulated maximum likelihood approach similar to that summarized in Train (2009). The simulated maximum likelihood estimation approach has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. Since we use panel data, the likelihood function at the family level is computed for a *sequence* of choices from t_0 to t_2, 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. Since the estimation algorithm is similar to a standard approach, we describe the remainder of the details in Appendix B.

5 Choice Model Results

Table 7 presents the results of the choice model. Column 1 presents the results from the primary specification while columns 2-5 present the results from four robustness analyses.

[39] Y_k is taken as time-invariant and relates to the maximum of dependents covered between t_0 and t_2.

[40] Since the model is a 'realized' utility model in dollar units, we don't need an additional scale normalization for the estimated variances.

In the primary specification, the switching cost intercept η_0 is large in magnitude with values of $1,729 for single employees and of $2,480 for employees who cover at least one dependent. An employee who enrolls in a flexible spending account (FSA) is estimated to have $551 lower switching costs than one who does not. The results show a small and negative relationship between income tier and switching costs. The coefficient describing the relationship between switching costs and being a manager (higher-level / white-collar employee) is positive but statistically insignificant while the coefficient linking switching costs to quantitative aptitude is near zero and insignificant. Employees (or their dependents) who have chronic medical conditions or a salient change in medical expenditures are estimated to have slightly higher switching costs than those without. This goes against our hypothesis that these employees would have lower switching costs because they pay closer attention to their insurance product. However, these employees are predominantly high-cost consumers who also may be unwilling to switch to a plan that they have no experience using.[41]

Since our estimates link switching costs to multiple dimensions of observable heterogeneity, we also present the population mean and variance of switching costs implied by our estimates. The mean total switching cost per employee is $2,032 with a population standard deviation of $446. This implies that, on average, when an employee has a previously chosen plan as their current default option, he forgoes up to $2,032 in expected savings from an alternative option to remain in the default plan. These results can also be viewed in light of the potential underlying sources of switching costs discussed in section 7. For example, note that the base of family switching costs is approximately 1.4 times larger than the base of individual switching costs despite having roughly 3 times the money at stake in the health insurance decision. This suggests that a pure inattention model with probabilistic re-optimization is not the primary basis for the estimated switching costs, since in this case switching costs would reflect the entire change in money at stake.

As the CARA coefficients presented in table 7 are difficult to interpret, we follow Cohen and Einav (2007) and analyze these estimates in a more intuitive manner in table 8. The table presents the value X that would make an individual with our estimated risk preferences indifferent between inaction and accepting a gamble with a 50% chance of gaining $100 and a 50% chance of losing $\$X$.[42] Thus, a risk neutral individual will have $X = \$100$ while an infinitely risk averse individual will have X close to zero. The top section of the table presents the results for the primary specification. X is $94.6 for the mean / median individual, implying a moderate amount of risk aversion relative to other results in the literature, which we present at the bottom of the table. X is $92.2 for the 95th percentile of γ and $91.8 for the 99th, so preferences don't exhibit substantial heterogeneity in the context of the litereature. Finally, our estimates in Table 7 reveal that the mean of the distribution of γ is slightly increasing in age and income, though neither effect has a large impact on the interpretations in table 8.[43]

[41] This kind of learning is not explicitly modeled and is embedded in the switching cost estimates. Farrell and Klemperer (2007) cite learning about alternative products as one potential underlying source of switching costs.

[42] These figures are computed for an individual with mean age and mean income.

[43] The positive relationship between income and risk aversion may reflect that (i) higher income employees have a heuristic that makes them more likely to select higher coverage and (ii) we don't estimate heterogeneity in plan intercepts with respect to income.

The results in table 7 also indicate that, above and beyond out-of-pocket expenditure risk, there is a strong distaste for PPO_{1200}. The distribution of the random coefficient δ for single employees has a mean of $\$-2,912$ with a standard deviation of \$843. Moreover, this coefficient internalizes the HSA match for first time enrollees of up to \$1,200, implying that the actual distaste for this plan is larger than the estimate indicates. This plan was, if anything, marketed more strongly to employees than the other options. The primary explanations for this distaste are (i) hassle costs from using the health savings account (ii) uncertainty surrounding how to use the health savings account for medical expenses and (iii) uncertainty about the retirement benefits of health savings accounts. We believe that decomposing the sources of this distaste is an interesting topic for future work.

Robustness. Table 7 presents results from four robustness specifications that provide insight into the sensitivity of the primary estimates with respect to core underlying assumptions. Column 2 studies a basic specification that estimates the model without conditioning the switching cost η and risk preference γ on observable demographics and linked choices. This sheds light on how the switching cost and risk preference estimates change when estimating these values for the whole population rather than as a function of observable heterogeneity.

Column 3 studies the impact of our cost model assumption that there is no moral hazard. To do this we necessarily make some simplifying assumptions: for a full structural treatment of moral hazard in health insurance utilization see, e.g., Cardon and Hendel (2001), Einav et al. (2011) or Kowalski (2011). We implement the moral hazard robustness check by adjusting the output of the cost model to reflect lower total utilization in the less comprehensive plans (and vice-versa). The intent is to show that, even when including price elasticities that are quite large relative to those found in the literature, the model output for switching costs and risk preferences does not change substantially. This analysis also sheds light on whether deviations in beliefs from F_{kjt}, e.g. from private information, have a marked impact on our results. Since this is a non-trivial exercise, we present the details of this analysis in Appendix C.

Column 4 studies the case where risk preference heterogeneity is log-normally distributed in order to determine sensitivity with respect to the normality assumption on γ. Column 5 follows Einav et al. (2011) and investigates the choice model without the family-plan-time specific idiosyncratic preference shock ϵ_{kjt}. As in their setting, there is a theoretical rationale for excluding this part of the model: the plans we study are vertically differentiated by financial characteristics but have no horizontal differentiation (except for HSA account linked to PPO_{1200} already modeled through δ). This robustness check allows us to assess the sensitivity of our parameters to the presence of ϵ and gain more insight into the empirical impact of this theoretically motivated restriction.[44]

Overall, the results from these alternative specifications suggest that our key parameter estimates and, consequently, the results from our information provision counterfactual analysis are robust to changes of the choice model's underlying assumptions. The population mean across Columns 2-5

[44] In our setting, ϵ for PPO_{1200} could be interpreted as time-varying preferences for the HSA option. For the other plans it could be a reduced form representation of deviation from the health expense expectations assumption.

ranges from \$1886 to \$2087 while the population variance ranges from \$286 to \$731. The specific coefficients for all model components are very similar to those in the primary specification for three of the four robustness analyses. The risk-preference robustness check (column 4) estimates differ somewhat: the mean level of switching costs is similar to that in the other specifications but the standard deviation is twice as large. This reflects the differing estimated coefficients on observable heterogeneity in this specification, which are larger in magnitude than our primary results. These differences likely arise from the choice implications of the wider tails of risk preferences inherent to the log-normal assumption, which we illustrate in the bottom section of table 8. We now turn to the issue of how reductions to the large switching costs that we find impact consumer choices, adverse selection, and welfare in our setting.

6 Information Provision: Switching Costs and Adverse Selection

In this study, consumers are enrolled in sub-optimal health plans over time (from the individual perspective) because of switching costs. After initially making informed decisions, consumers don't perfectly adjust their choices over time in response to changes to the market environment (e.g. prices) and their own health status. In this section, we use the results from the structural consumer choice analysis together with a model of information provision and health plan pricing to investigate the impact of a counterfactual policy that improves consumer choices by reducing switching costs. In each period, consumers are provided information that reduces their level of switching costs to some amount between their estimated switching cost level and zero. We study the welfare consequences of reduced switching costs in both (i) a 'naive' setting where the price of insurance does not change as a consequence of incremental selection and (ii) a 'sophisticated' setting where plan prices change to reflect the new risk profile of employees enrolled in the different options. While we describe this exercise as an information provision intervention, the following analysis can be applied to any policy or market change that leads to a reduction in switching costs. In section 7, we discuss specific information provision policies that have been proposed in the Affordable Care Act and employer provided insurance market, in the context of the underlying sources of switching costs that they impact.

Model of Information Provision and Plan Pricing. Formally, we model the impact of information provision by assuming that the improved consumer knowledge reduces switching costs to a fraction Z of the family-specific estimate η_k. Here, Z decreases as the counterfactual information provision policy becomes more effective. In the limit, as Z goes to zero, one could imagine a policy that leads to full re-optimization in each choice period.[45] To reflect the fact that only switching costs, plan prices, and past plan choices change in the counterfactual exercise from the current period choice perspective (health plan characteristics and health status are as observed) we restate the expected utility of family k for plan j at time t as an explicit function of these changing choice

[45] We assume that the information provision policy is costless, though the analysis could be performed where the policy has a cost that increases as Z declines.

factors:

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

We omit the dependence of utility on the other choice factors modeled in section 4 for notational simplicity, though we use continue to use all of these factors in the expected utility calculations. Consumers choose the plan j that maximizes their expected utility in each period t, subject to the preference estimates from our primary specification and the updated choice environment resulting from information provision. With endogenous plan pricing, these choices determine health plan costs, which in turn determine health plan premiums. As a result, the plan a family enrolls in in the environment with reduced switching costs depends both on the direct effect on their choice as well as the indirect effect that this reduction has on premiums resulting from the new profile of choices in the entire population. In theory, this collective externality on premiums from reduced switching costs could cause incremental advantageous selection, where the relative price of more comprehensive insurance decreases, or adverse selection, where this relative price increases.

In order to determine the impact of this externality in the information provision counterfactual, we model insurance plan pricing. Our model follows the pricing rule used by the firm during the time period studied, and is similar to plan pricing models used in the literature on the welfare consequences of adverse selection across a variety of contexts (see e.g. Cutler and Reber (1998) or Einav et al. (2010b) or Carlin and Town (2009)). The firm we study was self-insured for the PPO options in the choice model, implying that it has full control over the total premiums for each plan option as well as the subsidies employees receive toward those premiums. The total premium paid (by employer and employee), TP^y_{jt}, for each plan and year was set as the average plan cost for that plan's previous year's enrollees, plus an administrative markup, conditional on the dependent coverage tier denoted y:[46]

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

Here, $K^y_{j,t-1}$ refers to the population of families in plan j at time $t-1$ in coverage tier y. $PP_{kj,t-1}$ is the total *plan paid* in medical expenditures conditional on y and j at $t-1$. This total premium is the amount an employee in dependent category y enrolling in plan j would have to pay each year 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 conditional on the family's income tier, I_k.[47] Denote this subsidy $S(I_k)$. Building on these elements, the family-plan-time specific out-of-pocket premium P_{kjt} from the choice model is:

$$P_{kjt} = TP^y_{jt} - S(I_k)TP^y_{PPO_{1200}t}$$

[46] There are four coverage tiers (i) single (ii) employee plus spouse (iii) employee + child(ren) and (iv) employee plus spouse plus child(ren).

[47] The subsidy rates for the five income tiers ordered from poorest to wealthiest are .97, .93, .83, .71, and .64.

For PPO_{1200}, P_{kjt} is a fixed percentage of the total premium. For the other two PPO plan options employees pay the full marginal cost of the total premium relative to PPO_{1200}. Finally, since P_{kjt} depends on past cost information, we assume that P_{kj,t_0} equals the actual employee premium contributions set by the firm at t_0.[48] In section 7 we provide further motivation for our insurance pricing model in the context of the literature and discuss several alternative market models.

Welfare. We analyze welfare using a certainty equivalent approach that equates the expected utility for each potential health plan option, U_{kjt}, with a certain monetary payment Q. Formally, Q_{kjt} is determined for each family, plan, and time period by solving:

$$u(Q_{kjt}) = -\frac{1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)(W-Q_{kjt})} = U_{kjt}(P_{kjt}, Z\eta_k, \mathbf{1}_{kj,t-1})$$

The certainty equivalent loss Q_{kjt} makes a consumer indifferent between losing Q_{kjt} for sure and obtaining the risky payoff from enrolling in j. This welfare measure translates the expected utilities, which are subject to cardinal transformations, into values that can be interpreted in monetary terms.

An important issue in our setting is whether or not switching costs themselves should be incorporated into the welfare calculation as they change according to Z in the information provision exercise. It is natural to think that certain potential sources of our estimated switching costs should be excluded from the welfare calculation, while others imply a tangible social cost that should be incorporated. Since our empirical choice framework does not distinguish between the potential switching cost sources discussed in section 7, we study a range of welfare results spanning the case where switching costs aren't incorporated into the welfare calculation at all to the case where they are fully incorporated. Formally, we calculate the certainty equivalent loss as a function of the proportion of estimated switching costs that enter the welfare calculation, denoted κ:

$$u(Q_{kjt}^{\kappa}) = -\frac{1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)(W-Q_{kjt}^{\kappa})} = U_{kjt}(P_{kjt}, Z\eta_k, \mathbf{1}_{kj,t-1}, \kappa Z\eta_k)$$

As κ decreases from 1 to 0, the proportion of switching costs factored into the certainty equivalent for a non-incumbent plan decreases to $\kappa Z\eta_k$ and that potential switch becomes more attractive from a welfare perspective. Within this context, we investigate the welfare consequences of each information provision policy, described by Z, for κ between 0 and 1.[49] To our knowledge, this is the first empirical structural analysis that accounts for this range of welfare possibilities for a phenomenon like switching costs.

Conditional on κ, the family welfare impact of an information provision policy leading to a switching cost reduction to $Z\eta_k$ is:

[48] Presumably these contributions were set with the expectation that total premiums for each plan would equal average cost, though maintaining this stance is not necessary to assess the impact of the information intervention in the environment studied unless the firm would have adjusted P_{kj,t_0} together with the policies determining Z.

[49] This analysis relates to the welfare foundations laid out in Bernheim and Rangel (2009), who study a framework where choices can be close to, but not completely reflect, fundamental underlying preferences. Our approach is agnostic since it incorporates the entire range of potential treatments.

$$\Delta CS^Z_{kjt} = Q^{\kappa}_{k,j_Z,t} - Q^{\kappa}_{kjt}$$

This is the difference in certainty equivalents, for a given family, between the health plan chosen after the policy intervention, denoted j_z, and the choice j made in the benchmark model, net of any additional switching costs incorporated.[50] Since total premiums relate directly to average costs, the total welfare change differs from the consumer welfare change only if the sum of employee contributions P_{kjt} differs under policy Z than in the benchmark model. This will happen mechanically in our model since the market environment and consumer choices change as a result of information provision and the subsidy rule is not tied directly to an overall proportion of $\Sigma_k P_{kjt}$. Given this, we compute the mean per-family welfare change as follows:

$$\Delta TS^Z_t = \frac{1}{\|K\|}\Sigma_k \ \Delta CS^Z_{kjt} + \frac{1}{\|K\|}\Sigma_k \ (P^Z_{kjt} - P_{kjt})$$

The distinction between consumers surplus and total surplus here depends only on this change in aggregate premiums paid and is not a substantive issue. If the employee contributions were required to add up to a portion of total premiums conditional on Z then consumer welfare is equivalent to overall welfare in our model.[51] Since aggregate premiums do not change as the enrollment profile does, the welfare change here results primarily from impact of risk preferences as consumers are matched to different plans. We report this welfare change as a percentage by dividing ΔTS^Z_t by three different metrics from the benchmark setting (i) the average employee premium paid in year t (ii) the average sum of premium and out-of-pocket medical expenditures at t and (iii) the average total certainty equivalent loss of the plans consumers enroll in at t. These calculations give a sense of the welfare consequences of the information intervention relative to important quantities: for a further discussion of potential welfare benchmarks in health insurance markets see Einav et al. (2010b).

Results: No Plan Re-Pricing. Before we investigate the interaction between a reduction in switching costs and adverse selection, we analyze the 'naive' case where information provision occurs but health plan premiums are held fixed as observed in the data. Consumers may switch to a new health plan as a result of lower switching costs, but this selection does not feed back into prices on the basis of new enrollee cost profiles. In this context, the policy intervention can only increase welfare since prices are by definition unchanged and the policy helps consumers make weakly better decisions relative to the benchmark case. This analysis presents a direct comparison to prior work that studies the impact of reduced choice frictions where (i) consumer choices don't impact firm costs or (ii) the feedback between choices and costs/prices is ignored.[52]

[50] Note that this family-level welfare impact will generically be non-zero since premiums will adjust as the result of the intervention Z even if a given family's choice does not.

[51] In this case, total premiums could be held constant moving to policy Z by taking the per-person difference in total premiums across two environments and adding or subtracting this term from P^Z_{kjt}.

[52] For prior work of this kind on insurance markets see e.g. Kling et al. (2011) and Abaluck and Gruber (2011). More broadly, there is some work studying the equilibrium consequences of consumer choice inadequacy in the industrial organization literature, though these analyses do not study a context where choices directly impact costs as they do

Figure 3 presents market share and average cost results for the partial equilibrium information provision simulation as a function of Z. Results are presented for years t_1 and t_2 for switching cost levels ranging from 0 to η. The top panel reveals that, as Z decreases towards 0, consumers adjust more readily to the observed price changes and are more likely to enroll in PPO_{500} (recall relative prices for PPO_{250} went up in t_1). For the case where $Z = \frac{1}{4}$, and information provision removes most but not all switching costs, 913 employees enroll in PPO_{500} at t_1, a 44% increase over the benchmark model with full switching costs where 639 consumers choose that plan. For the cases of $Z = \frac{1}{2}$ and $Z = 0$, t_1 enrollments in PPO_{500} are 780 (21% increase) and $1,052$ (65% increase) respectively. Moving forward to t_2, for $Z = \frac{1}{4}$, there are $1,010$ enrollees in PPO_{500} relative to 702 in the benchmark case (a 44% increase). Almost all of the switchers towards PPO_{500} would have continued enrollment in PPO_{250} in the benchmark case. The figure also reveals that switching to and from PPO_{1200} as the result of information provision is limited, due to the horizontal differentiation δ resulting from the health savings account and linked features. The lower panel in figure 3 reveals that these different enrollment patterns over time imply different plan average costs. For the family coverage tier average costs for PPO_{250} increase relative to those in PPO_{500} as enrollment in the former declines relative to the latter, implying that the people who switch out of PPO_{250} into PPO_{500} are healthier than those who do not switch. This result is similar for other coverage tiers and suggests that, in the primary analysis with endogenous plan re-pricing, PPO_{250} premiums will become more expensive, potentially leading to even more selection against that plan in subsequent periods.

Table 9 presents the welfare impact of moving from the benchmark environment with full switching costs to the case where $Z = \frac{1}{4}$. At t_2, the mean per employee certainty equivalent increase is \$114. For those who switch plans in the counterfactual environment relative to the benchmark case, the mean benefit is \$196 (for those who do not switch, the change is zero by definition). The policy intervention improves welfare by 5.8% of total employee premium contributions, and by 2.5% relative to the total employee spending benchmark (the table also reports the % changes for the certainty equivalent at stake benchmark). The analagous percentage changes for those who switch plans are 10.0% and 4.4% relative to these respective metrics. These numbers are similar to, but slightly larger tham, the impact of the information intervention during t_1. The positive welfare impact of improved individual-level choices in the environment without plan re-pricing is similar to prior results in the empirical literature on choice inadequacy, but stands in stark contrast to the negative welfare results in our analysis with endogenous plan re-pricing, which we now turn to.

Results: Endogenous Plan Re-Pricing. With endogenous plan re-pricing, premiums change as consumers switch plans due to the policy intervention. It is possible that even a small change to the profile of choices without plan re-pricing will map to a large change in premiums and choices with endogenous re-pricing. A small enrollment change under no re-pricing could imply a change in premiums that leads to further incremental switching, in turn leading to further enrollment changes

when adverse selection is a concern (see e.g. Gabaix and Laibson (2006) or Ellison (2006)).

and an unraveling process that continues until it reaches a new fixed point between enrollment and premiums. The link between choices, costs, and prices could lead to more or less adverse selection over time in the presence of reduced switching costs. For example, in our empirical model switching costs are estimated conditional on multiple dimensions of observable heterogeneity that are correlated with health status in a specific manner. In the spirit of Einav et al. (2010c) and Cutler et al. (2008), if estimated switching costs are sufficiently negatively correlated with expected expenditures then healthy people who initially sign up for PPO_{250} would be less likely to switch from that plan over time conditional on the value of switching, leading to reduced adverse selection. More broadly, the nature of incremental selection depends on both the stochastic process governing health status and the choice environment when consumers make initial choices, in the absence of switching costs. The results from our analysis with no plan re-pricing suggest that we will find increased adverse selection as a result of reduced switching costs since, across the range of Z, we find that the average cost difference between PPO_{250} and PPO_{500} increases.

For each Z, we study the evolution of choices, prices, and welfare from year t_0 to t_6, four years beyond the end of our data.[53] The top panel in Figure 4 presents the time path of plan market shares for PPO_{250} and PPO_{500} and the two cases of (i) $Z = 1$ (full switching costs) and (ii) $Z = \frac{1}{4}$. The impact of the policy intervention on the market share of PPO_{250} relative to PPO_{500} is noticeable: reduced switching costs decrease t_6 enrollment in PPO_{250} from 744 to 385 and increase enrollment in PPO_{500} from 647 to 1134. This indicates that the improved choices over time substantially increase incremental adverse selection, to the point where PPO_{250} is almost eliminated from the market due to high premiums caused by the sick profile of enrollees (this kind of insurance market unraveling is known as a 'death spiral', see e.g. Cutler and Reber (1998)). This enrollment gap is also large for years t_1 to t_5. Relative to the no re-pricing case, in t_1 and t_2 PPO_{250} has much lower enrollment after the policy intervention, revealing the large impact of endogenous re-pricing. The bottom panel of Figure 4 reveals the substantial and increasing family tier average cost differential between these two plans over time. The average cost (and total premium) of PPO_{250} increases relative to PPO_{500} under information provision for all years from t_1 to t_6, with a maximum relative change of $\$4,619$. This pattern is similar for the other coverage tiers and indicates significant incremental adverse selection as a result of reduced switching costs. Enrollment in PPO_{1200} does not change markedly relative to the benchmark case, due to the substantial horizontal differentiation.

Figure 5 shows market shares and average costs for $PP0_{500}$ and PPO_{250} in years t_1, t_2, and t_4 as a function of Z. We report results for values of Z equal to 0,.25,.5, and 1. It is clear from the top panel that as Z decreases from 1, and information provision becomes more effective, enrollment in PPO_{250} declines at the expense of enrollment in PPO_{500}. The bottom panel reveals that as Z decreases and enrollment increases, the relative average costs (and total premiums) of PPO_{250} increase, implying that as information provision becomes more effective we find higher incremental adverse selection.

[53] From t_3 to t_6 we assume that the demographics and health status for the sample are the same as observed in t_2. This implies that the analysis for these time periods reflects the long-run impact of reduced switching costs on prices and selection.

Table 10 presents a detailed analysis of the welfare impact of an information provision policy that reduces switching costs from η to $.25\eta$ ($Z = .25$). For this table and table 11, we assume that $\kappa = 0$, implying that the reduction in switching costs from the policy intervention does not enter the welfare calculation.[54] We calculate this impact for the population overall as well as for select groups of interest. For the entire population, the information provision policy to *improve* choices has a *negative* welfare impact in each year and overall. The mean per employee per year certainty equivalent welfare loss is \$115, implying an average per person welfare loss of \$690 over the time period studied. This translates to a 7.7% loss using average employee premium contributions as a benchmark. Table 10 also reveals that the information provision policy has substantial distributional consequences. Employees who switch plans as a result of the intervention have an average welfare gain of \$186 per employee per year (12.4% of total premiums). Those who do not switch plans experience a mean per employee per year loss of \$442 (−29.4%). The welfare impact on employees that do not switch plans comes entirely from the changes to their plan prices resulting from incremental selection from those that do switch. This is interesting to contrast with the results under no re-pricing, where non-switchers have zero welfare loss by necessity.[55] The table also shows that high-expense employees experience a small welfare improvement from the intervention equivalent to \$62 (4.1%) per employee per year versus a \$137 (−9.1%) per employee per year loss for all other employees.[56] Single employees lose an average of \$319 (−21.3%) per employee per year versus a \$61 gain for employees with at least one dependent. Finally, lower income employees (making less than \$72,000 per year) lose an average of \$200 per employee per year as a result of the intervention, versus no change in welfare for high income employees. While these specific group effects may differ in other contexts, it is likely that any policy that substantially improves (or hinders) choices in similar health insurance markets will have non-trivial distributional consequences.

Table 11 broadens the welfare analysis to include a spectrum of information provision counterfactuals, including the cases of Z equal to 0.75, 0.5, 0.25, and 0. We also study the welfare loss of adverse selection in our observed setting relative to the conditional first-best outcome. In our setting, the welfare loss comes from (i) having risk-averse individuals enrolled in less comprehensive insurance and (when $\kappa > 0$) (ii) switching costs. Thus, the first-best conditional on the set of health plans offered is to have all employees enrolled in PPO_{250} in every time period. The table presents the average per employee per year certainty equivalent change for each of these counterfactual scenarios relative to the baseline (the analogous figures to the last column in Table 10), with percentage changes reported relative to average employee premium contributions. The welfare loss from adverse selection in the baseline relative to the first best is \$123 (−8.2%) per employee per year.[57] This welfare loss becomes worse as Z declines from 1 to 0.25. For Z equal to .75, .5,

[54]Thus, the welfare impact here reflects purely the welfare difference from the choices made and does not incorporate the assumed reduction in switching costs as a tangible cost.

[55]Here, total premiums increase in both PPO_{250} and PPO_{500}, implying that every family that remains in the same plan loses money relative to case with full switching costs.

[56]High-expense is defined as spending more than \$15,000 for a single employee, \$25,000 for a family of size two, and \$32,000 for a larger family (this covers approximately 10% of the population).

[57]This can be compared directly to the numbers found in the literature on the welfare consequences of adverse selection. Cutler and Reber (1998) find that the welfare loss from adverse selection in their environment is between

and .25 the incremental losses relative to our baseline are \$41 (-2.7%), \$73 (-4.9%), and \$115 (-7.7%) respectively. For $Z = 0$, this loss is \$107 (-7.1%),a slightly smaller loss than that when $Z = 0.25$.[58]. The table also presents the welfare impact for the select population groups studied in table 10. Employees who switch plans as a result of the intervention have a substantial welfare gain equivalent to \$1,017 (68%) per person when $Z = 0.75$, which decreases to a \$118 (7.9%) when $Z = 0$. For those who don't switch, the welfare loss ranges from \$249 (-16.6%) in the former case to \$382 (-25.4%) in the latter. The results for employee groups based on coverage status, income status, and expense level are similar in direction to those presented for $Z = .25$ in table 10, though there is some variation in their relative magnitudes. The distributional impacts of moving to the first-best are substantial, and reflect the pooling of the entire population into $PP0_{250}$. The sick and high-income employees who are more likely to be in this plan in the baseline benefit from pooling at the expense of the healthier low-income employees more likely to be moved into this plan.

Table 12 expands the analysis further to allow for the possibility that some proportion of estimated, and subsequently reduced, switching costs be incorporated into the welfare calculation. The top panel studies the profile of switching costs incurred under different Z over time. The mean per employee per year switching costs incurred is \$185 for the baseline case and \$188 for the case where $Z = 0.75$ (note that this calculation covers the entire population, not switchers only). This means that the product of the number of switchers and switching costs is slightly larger as switching costs are reduced from the baseline, implying that the increased number of switchers makes up for the decrease in switching costs per switcher. For $Z = .5$ and $Z = .25$ the mean per employee per year incurred switching costs are \$142 and \$83 respectively, while $Z = 0$ by definition implies no switching costs. These patterns are similar for each individual year, though the overall level of switching costs incurred is higher in t_1 and t_2 than in later years as the market moves towards a steady state.

The bottom panel of table 12 presents the overall welfare impact of the range of policies described by Z, as a function of κ, together with the extent to which switching costs impact this calculation in each case.[59] Results are presented for κ equal to 0 (no switching costs in welfare calculation), .25, .5, and 1 (all included). The results for $\kappa = 0$ restate the results from table 11. When $\kappa = 0.25$, the welfare change caused by the policy intervention is \$90 (-6.0%) for $Z = 0.25$, compared to the \$115 loss when $\kappa = 0$. In the extreme case when all switching costs count in the welfare calculation ($\kappa = 1$), the $Z = 0.25$ policy leads to a \$13 (-0.9%) loss. More broadly, the table reveals that the impact of including a higher proportion of switching costs in the welfare calculation is larger the more switching costs are reduced by the intervention. For example, the welfare impact of $Z = 0.75$

2-4% of total baseline spending. In our setting, when renormalized by this metric, we find a welfare loss of -2.9%, right in the middle of this range. Recent work by Carlin and Town (2009), Bundorf et al. (2010), Einav et al. (2010b), and Einav et al. (2011) find results that are in this ballpark in a variety of empirical settings.

[58] This non-monotonicity of the welfare loss in Z near 0 arises because, when switching costs are entirely removed, premiums and enrollment oscillate towards a stable equilibrium rather than move monotonically to that steady state. This results from the specific features of the distribution of health risk and evolution of prices and is characteristic of the intricacies of equilibria with adverse selection (see e.g. Rothschild and Stiglitz (1976))

[59] Welfare relevant switching costs equal κ multiplied by the average per employee per year population switching costs incurred for each policy intervention (the bottom row in the top panel).

relative to the baseline is close to constant as a function of κ, while the more effective $Z = 0.25$ intervention has a wider range of potential welfare impacts as a function of κ. Importantly, the welfare impact of information provision is negative across the range of κ for all Z except $Z = 0$. This suggests that regardless of the stance taken on what the sources of the estimated switching costs are, the counterfactual reduction in switching costs will lead to increased adverse selection and lower welfare in general.[60]

7 Discussion

Sources of Switching Costs. In the empirical framework switching costs enter the expected utility function as a fixed monetary preference for the incumbent plan relative to alternative options. However, it is likely that the substantial switching costs documented in this work arise from a variety of sources tied to different underlying microeconomic fundamentals. If alternative sources of switching costs are prominent this could have implications for (i) the choice model estimates (ii) what policies are actually effective in reducing switching costs and (iii) how switching costs interact with adverse selection. We now present some prominent alternative explanations and argue that differentiating between these explanations would not markedly changes our results on the impact of switching costs and their interaction with adverse selection.

Transactions Costs. Transactions costs are the most commonly cited potential source of switching costs in the theoretical and empirical literatures (see e.g. Farrell and Klemperer (2007) for the former and Shum (2004) or Dube et al. (2010) for the latter). Our choice framework follows the prior empirical literature and models switching costs as if they arise completely from transaction costs. In our setting, transaction costs could result from, e.g., the time and hassle costs of (i) researching alternative health plan options or (ii) actually switching plans. It is clear that transaction costs have the potential to be quite large in our setting. Health plans are complicated objects to choose between, often described in long documents with substantial legalese, while the process of actually making a choice can require multiple non-trivial and costly actions.

Learning Costs. Learning costs are another potential source of switching costs (see e.g. Klemperer (1995)). Learning costs occur when, after purchasing a product, a consumer has to put in time and effort to learn about how to use its various features. Learning costs imply that once you have purchased a product, you don't incur this source of switching costs when moving back to that product. Here, this is likely only a small component of the switching costs we estimate since the three PPO plan options have the exact same network of providers and the only real product-specific learning that could occur relates to the Health Savings Account (HSA) option in PPO_{1200}.

[60]For $Z = 0$, the welfare impact of the intervention is negative for κ of 0,0.25, and 0.5 but positive for $\kappa = 1$. This means when all switching costs are a tangible welfare component and switching costs are entirely eradicated, the information intervention increases welfare. In this case, there is still increased adverse selection from the intervention, but the positive welfare effect from the substantial reduction in switching costs paid outweights the negative impact of adverse selection.

Product Compatibility. Learning costs and product compatibility, another potential source of switching costs, would both be important issues in cases where consumers had to switch medical providers and disrupt continuity of care when switching health plan. Compatibility costs occur when product-specific investments are made that increase the value of a product relative to an otherwise similar product (see e.g. Klemperer (1995)). In our setting, since the PPO plan options we study have the same network of providers, there are no compatibility issues because relationship-specific investments can be transferred seamlessly across these choices. This implies that the switching costs we estimate are a lower bound relative to switching costs in a broader plan setting where a consumer has to switch providers to switch health plan. We believe that switching costs from provider lock-in are an important issue to study in future work.

Fixed Re-Optimization Cost. This model and prior work studies switching costs in an environment where the models assume that consumers have perfect knowledge of product characteristics. In the transaction costs section above, we noted the research costs are a potentially important source of switching costs, implying that there may be imperfect knowledge of plan options. In this case, a more sophisticated model could treat the decision process as a two-stage process where (i) consumers pay some transaction cost to re-optimize and (ii) change plans after acquiring information. This is different than the transaction costs explanation above because ex ante beliefs about product characteristics are important to the decision to re-optimize: a consumer could have small transaction / re-optimization costs but beliefs that they have little to gain from acquiring information. In our setting, where prices changed substantially from t_0 to t_1, consumer beliefs may have been centered around status-quo prices, implying that a small transaction cost could dissuade them from re-optimizing and in turn lead to a large switching cost estimate. In this case, our switching cost estimates can be interpreted as money that consumers leave on the table as a result of this decision process. Note that since there is low switching for t_2 as well, under this model consumers in our environment also do not learn about the potential to gain a lot of money from switching over time, implying incorrect beliefs can be quite persistent.

Inertial & Psychological Costs. There is a substantial empirical literature studying consumer inertia and psychological switching costs (see e.g. Madrian and Shea (2001), Samuelson and Zeckhauser (1988), or Thaler and Sunstein (2008)). This explanation is also related to the empirical literature on brand loyalty in marketing (see e.g. Dube et al. (2009)). Following Klemperer (1995) we define this source of switching costs as those that arise above and beyond any clearly identifiable rational phenomenon.[61]

These different potential sources could have different implications for how switching costs inter-

[61] The empirical literature on inertia overlaps with the literature on switching costs in that some of the phenomena we describe in this discussion are also sources of inertia in the literature. Thus, our definition of inertia here is narrower than that in the literature, while our definition of switching costs contains all sources of inertia in the literature.

act with adverse selection. For example, if inertia or the fixed re-optimization cost are prominent sources then people who switch on the margin would come more diffusely from the distribution of the marginal benefit from switching plan. If transaction costs are prominent, then those with the most to gain from switching will be more likely to do so than others who benefit, but by less. However, in either case, as switching costs are reduced in the counterfactual exercise we would be moving from the environment we observe with substantial choice persistence to one where consumers are making almost entirely active choices. The environment with reduced switching costs would have similarly increased adverse selection in either scenario, leading to similar welfare implications for the policy intervention. Additionally, it is unlikely that the model of switching costs seriously biases other parameter estimates (such as risk preferences) because those parameters are identified separately from choice persistence (in any form) based on choices made in the active choice year at t_0, which will not change regardless of the specific model of switching costs.

Of course, these arguments do not determine whether switching costs themselves are a net social cost: transaction costs and learning costs should be treated as social costs while the implied impact of the fixed re-optimization cost (above the cost itself) and inertial costs should likely not be included in the welfare calculus.[62] We address this issue agnostically in our results section by presenting the welfare implications of the policy intervention over the possible situations ranging from full inclusion to no inclusion. We believe that separately identifying these different underlying sources is fertile ground for future work.

Policies That Reduce Switching Costs. There is a general consensus in the policy debate on the design and regulation of health insurance markets that helping consumers make the best plan choices possible is unequivocally the right course of action, regardless of the specifics of the environment.[63] The recently passed Affordable Care Act (ACA), sets forth guidelines for insurance exchanges that mandate clear and simple information provision for plan options (see e.g. Kaiser Family Foundation (2010b). In the employer insurance market many employers, including the one we study here, have implemented web-based decision tools to try and clearly explain benefits to employees (Kaiser Family Foundation (2010a)). Both of these policies have the potential to substantially reduce switching costs stemming from either transaction costs or a fixed re-optimization cost. Other policies, such as targeted default options or no default option could reduce the impact of switching costs in the spirit of our counterfactual if inertia is a prominent source. Simple individual-specific plan recommendations or streamlining the online plan choice process are additional ways to reduce multiple switching costs factors. Going forward, these policies, along with a multitude of other potential interventions, will provide decision-makers with a variety of tools to reduce the

[62]If a large part of the switching costs we estimate come from the combination of a small re-optimization cost and specific ex ante beliefs about price and health changes, then we believe that the actual re-optimization cost should count as a welfare loss but the full estimated switching cost should not.

[63]Advocates of the health insurance exchanges proposed in the Affordable Care Act (ACA) stress that a well-informed, optimizing consumer base is essential to realize the benefits from insurer competition in these markets (see e.g. Enthoven et al. (2001)). In the employer-sponsored insurance market there has been a similar emphasis on providing consumers increased plan options and the capabilities to choose between them Kaiser Family Foundation (2010a).

level of switching costs and impact choice in insurance markets.

Insurance Pricing. The two primary parts of the insurance pricing model are (i) average cost-plus pricing and (ii) a lump sum subsidy for any plan chosen. The former arises naturally as the product of competition while the latter is a policy designed to encourage consumers to internalize price signals and choose plans more efficiently. However, an environment where consumers fully internalize marginal prices also leads to increased adverse selection since, in general, a smaller and more expensive profile of consumers will select comprehensive insurance as its relative price increases.[64] In our context, this implies that as switching costs are reduced, consumers change plans, and total premiums change, the employee premium contribution P_{kjt} will change by the same amount that the total premium TP^y_{jt} does. This leads to higher incremental adverse selection relative to environments where subsidies increase as a function of the total premium.

These features are core components of the managed competition paradigm that motivates the health insurance exchanges proposed in the ACA (see e.g. Enthoven et al. (2001) and Kaiser Family Foundation (2010b)) and are also common in the current employer sponsored insurance market.[65] Cutler and Reber (1998) propose a risk-adjustment mechanism (incorporated into the ACA) to layer on top of average cost pricing and lump sum subsidies that will mitigate adverse selection in the market. Under their framework, employers (or the exchange regulator) can arrange transfers between insurers to reflect differences in the ex ante risk profile of the consumers these plans enroll. These transfers link directly to the average cost differences between the plans such that, in an ideal setting, plans set premiums that reflect consumer valuations for them netting out expected health expenditures. Consumers still face the full marginal premium difference in their contributions, but now these differences are smaller than those with risk-adjustment because of the transfers. In our setting, incorporating risk-adjustment into premium setting would lead to a welfare impact bounded between the first best and the results we found for each information provision setting. The degree to which welfare would improve towards the first-best depends both on (i) the effectiveness of the risk-adjustment scheme and (ii) the distribution of population preferences.

Finally, it is also interesting to consider how firms competing in a free market would set prices in response to the switching costs we estimate. There is an extensive literature surveyed in Farrell and Klemperer (2007) that considers how firms compete in environments with switching costs. This work shows that competition in markets with switching costs, even in the absence of adverse selection, can be quite complex to analyze. For example, Cabral (2008) and Dube et al. (2009) both investigate environments where switching costs can actually increase competitiveness through increased up-front competition for consumers. We are unaware of theoretical work studying unregulated competition in the presence of switching costs and adverse selection and feel future work in this area could be beneficial.

[64] We note that this is true as long as additional dimensions of heterogeneity, such as preferences, are not highly negatively correlated with expected health expenditures (see e.g. Einav et al. (2010c) or Cutler et al. (2008)).

[65] Under an almost identical pricing regime, Cutler and Reber (1998) study the trade-off between competition and adverse selection resulting from a shift towards lump sum subsidies for the health plans offered at Harvard University. The authors also discuss multiple empirical cases with similar pricing environments with evidence of adverse selection.

8 Conclusion

In this paper, we leverage a unique natural experiment in combination with individual-level choice and medical claims data to cleanly identify switching costs in a health insurance market. Several model-free preliminary analyses reveal that switching costs have a substantial impact on health plan enrollment as the choice environment evolves over time. We estimate a choice model of decision-making under uncertainty to quantify switching costs, ex ante health risk, and risk preferences and use these outputs to study the impact of counterfactual policies that reduce switching costs in a context with adverse selection.

We find evidence of high and heterogeneous switching costs: in the base specification the mean population level is just over $2,000 with a population standard deviation near $500. In a counterfactual setting where we hold insurance prices fixed as observed in the data, reducing switching costs to one-quarter of the estimated values improves consumer welfare by approximately 5% of total premiums paid. However, in the primary counterfactual analysis where insurance prices endogenously re-adjust to reflect the new enrollment patterns, the same reduction in switching costs exacerbates adverse selection and *reduces* welfare by 7.7% of premiums paid. We show that welfare is decreasing as the intervention to reduce switching costs becomes more effective and that our main result is robust to different welfare interpretations of switching costs arising from the different potential underlying sources. This result can also be vieweed in the context of the welfare loss from adverse selection in our observed environment relative to the first-best. We show that the welfare loss from adverse selection in the observed environment is 8.2%, implying that the policy to reduce switching costs effectively doubles the welfare loss from adverse selection. Finally, the results reveal that there are substantial distributional consequences from the reduction in switching costs, in addition to the overall efficiency loss.

Since improving choice adequacy could have desirable effects not considered here (e.g. enhancing the efficiency of product offerings), and this is one specific setting, we don't take the position that keeping consumers uninformed is a good intentional policy tool. However, the analysis clearly reveals that policies to improve consumer choices can have unintended efficiency and distributional consequences in markets with adverse selection. Prior work on choice adequacy either applies to markets without the potential for adverse selection or ignores this possibility by focusing on the partial equilibrium impact of improved choices. Policymakers in the health insurance sector constantly stress the need to improve consumer decision-making capabilities and reduce consumer switching costs, but likely do so without simultaneously evaluating how other regulatory policies (e.g. insurance contract characteristics or subsidy schedules) should change to reflect society's goals as the level of decision-making improves. As the capability to reduce consumer switching costs and improve their choices increases, it will be important for regulators to re-assess and potentially revise these other market policies to avoid undesirable social outcomes. We believe that a formal analysis of interactions between policies to improve decision-making and other key market regulations in setting with potential adverse selection is a valuable direction for future research.

References

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

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

Bernheim, D. and A. Rangel, "Beyond Revealed Preference: Choice-Theoretic Fuondations for Behavioral Welfare Economics," *Quarterly Journal of Economics*, 2009, *124* (1), 146–163.

Bundorf, K., J. Levin, and N. Mahoney, "Pricing and Welfare in Health Plan Choice," April 2010. Stanford University, Working Paper.

Cabral, L., "Switching Costs and Equilibrium Prices," August 2008. New York University Working Paper.

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

Carlin, C. and R. Town, "Adverse Selection, Welfare, and Optimal Pricing of Employer Sponsored Health Plans," 2009. University of Minnesota Working Paper.

Chandra, A., J. Gruber, and R. McKnight, "Patient Cost-Sharing and Hospitalization Offsets in the Elderly," *American Economic Review*, 2010, *100* (1), 193–213.

Chiappori, P., A. Gandhi, B. Salanie, and F. Salanie, "Identifying Preferences under Risk from Discrete Choices," December 2008. Columbia University Working Paper.

_ **and B. Salanie**, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 2000, *108*, 56–78.

Cohen, A. and L. Einav, "Estimating Risk Preferences from Deductible Choice," *The American Economic Review*, 2007, *97* (3), 745–788.

Crawford, G., N. Tosini, and K. Waehrer, "The Impact of Rollover Contracts on Switching Costs in the UK Voice Market: Evidence from Disaggregate Customer Billing Data," July 2011. University of Warwick, Working Paper.

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

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

_ **, B. Lincoln, and R. Zeckhauser**, "Selection Stories: Understanding Movement Across Health Plans," July 2009. NBER Working Paper No. 15164.

DeNavas-Walk, C., B. Proctor, and J. Smith, "Income, Poverty, and Health Insurance Coverage in the United States: 2009," *U.S. Census Bureau*, 2010.

Dube, J., G. Hitsch, and P. Rossi, "Do Switching Costs Make Markets More Competitive," *Journal of Marketing Research*, 2009, *46*, 435–445.

_ , _ , and _ , “State Dependence and Alternative Explanations for Consumer Inertia,” *The RAND Journal of Economics*, 2010, *41* (3), 417–445.

_ , _ , _ , and M. Vitorino, “Category Pricing with State Dependent Utility,” *Marketing Science*, 2008, *27* (3), 417–429.

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

_ , _ , and M. Cullen, “Estimating Welfare in Insurance Markets Using Variation in Prices,” *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.

_ , _ , and P. Schrimpf, “Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market,” *Econometrica*, 2010, *78* (3), 1031–1092.

_ , _ , S. Ryan, P. Schrimpf, and M. Cullen, “Selection on Moral Hazard in Health Insurance,” April 2011. Stanford University Working Paper.

Ellison, G., “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Enthoven, A., A. Garber, and S. Singer, “Near-Universal Coverage Through Health Plan Competition: An Insurance Exchange Approach,” *Covering America: Real Remedies for the Uninsured, Editors J. Meyer and E. Wicks*, 2001, *Washington DC.*

Ericson, K., “Market Design When Firms Interact with Inertial Consumers: Evidence from Medicare Part D,” November 2010. Harvard, Working Paper.

Fang, H., M. Keane, and D. Silverman, “Sources of Advantageous Selection: Evidence from the Medigap Insurance Market,” *Journal of Political Economy*, 2008, *116* (2), 303–350.

Farrell, J. and P. Klemperer, “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” *Handbook of Industrial Organization Chapter 31 Editors M. Armstrong and R. Porter*, 2007, pp. 1967–2072.

Finkelstein, A. and J. Poterba, “Testing for Adverse Selection with Unused Observables,” 2006. NBER Working Paper 12112.

Gabaix, X. and D. Laibson, “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gertner, R., “Game Shows and Economic Behavior: Risk-Taking on ‘Card Sharks’,” *Quarterly Journal of Economics*, 1993, *108* (2).

Goettler, R. and K. Clay, “Tariff Choice with Consumer Learning and Switching Costs,” October 2010. University of Chicago Booth School of Business, Working Paper.

Handel, B., “Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts,” January 2010. Northwestern University Working Paper.

_ , I. Hendel, and M. Whinston, “Better Off Without You?: Hassle Costs and Foregone Benefits in Flexible Spending Accounts,” September 2011. Northwestern University Working Paper.

Holt, C. and S. Laury, "Risk Aversion and Incentive Effects," *American Economic Review*, 2002, *92* (5).

Kaiser Family Foundation, "Employer Health Benefits: 2010 Annual Survey," http://ehbs.kff.org/pdf/2010/8085.pdf 2010.

_ , "Focus on Health Reform: Summary of New Health Reform Law," http://www.kff.org/healthreform/upload/8061.pdf 2010.

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

Kling, J.R., S. Mullainathan, E. Shafir, L Vermeulen, and M.V. Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," June 2011. Princeton University, Working Paper.

Kowalski, A., "Estimating the Tradeoff Between Risk Protection and Moral Hazard with a Nonlinear Budget Set Model of Health Insurance," April 2011. Yale Department of Economics working paper.

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

McManus, M., S. Berman, T. McInerny, and S. Tang, "Weighing the Risks of Consumers Driven Health Plans for Families," *Pediatrics*, 2006, *117*, 1420.

Newhouse, J., *Lessons from the RAND Health Insurance Experiment*, Harvard University Press, 1993.

Reed, M., V. Fung, M. Price, R. Brand, N. Benedetti, S. Derose, J. Newhouse, and J. Hsu, "High-Deductible Insurance Plans: Efforts to Sharpen a Blunt Instrument," *Health Affairs*, 2009, *28* (4), 1145–1154.

Rothschild, M. and J. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 1976, *90*, 630–649.

Samuelson, W. and R. Zeckhauser, "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1988, *1*, 7–59.

Shum, M., "Does Advertising Overcome Brand Loyalty? Evidence From Breakfast Cereals," *Journal of Economics and Management Strategy*, 2004, *13*, 241–272.

Strombom, B., C. Buchmueller, and J. Feldstein, "Switching costs, price sensitivity, and health plan choice," *Journal of Health Economics*, 2002, *21*, 89–116.

Sydnor, J., "Over(insuring) Modest Risks," *American Economic Journal: Applied Economics*, 2010, *2* (4), 177–199.

Thaler, R. and C. Sunstein, *Nudge*, Yale University Press, 2008.

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

Appendix A: Cost Model Setup and Estimation

This appendix describes the details of the cost model, which is summarized at a high-level in section 4. The output of this model, F_{kjt}, is a family-plan-time specific distribution of predicted out-of-pocket expenditures for the upcoming year. This distribution is an important input into the choice model, where it enters as a family's predictions of its out-of-pocket expenses at the time of plan choice, for each plan option. We predict this distribution in a sophisticated manner that incorporates (i) past diagnostic information (ICD-9 codes) (ii) the Johns Hopkins ACG predictive medical software package (iii) a novel non-parametric model linking modeled health risk to total medical expenditures using observed cost data and (iv) a detailed division of medical claims and health plan characteristics to precisely map total medical expenditures to out-of-pocket expenses.[66] The unique level of precision we gain from the cost model leads to more credible estimates of the choice parameters of primary interest (e.g. switching costs).

In order to most precisely predict expenses, we categorize the universe of total medical claims into four mutually exclusive and exhaustive subdivisions of claims using the claims data. These categories are (i) hospital and physician (ii) pharmacy (iii) mental health and (iv) physician office visit. We divide claims into these four specific 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}.

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:

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

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 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:

$$\begin{aligned} \text{Pharmacy:} \quad & \lambda_{it} \\ \text{Hospital / Physician (Non-OV):} \quad & \theta_{it} \\ \text{Physician Office Visit:} \quad & \theta_{it} \\ \text{Mental Health:} \quad & C_{MH,i,t-1} \end{aligned}$$

For pharmacy claims, individuals are grouped into cells based on the predicted future mean phar-

[66] Features (iii) and (iv) are methodological advances. 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 (2009)) in a structural setup. Einav et al. (2011) use this type of framework in ongoing work.

macy 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 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}:

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

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').[67] 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 $\rho_{z_\theta}(d, d')$:

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

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:

$$\Omega_j : C \rightarrow OOP_j$$

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

[67] 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.

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 6 shows the results of one validation exercise for PPO_{250}. The top panel reveals that actual employee out-of-pocket spending amounts are quite close to those predicted by Ω_j, indicating the precision of this mapping. The bottom panel repeats this mapping when we add out of network expenses as a fifth category. The output in this case is similar to that in the top panel without this category, implying that including this category would not markedly change the cost model results.

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. Figure 7 outlines the primary components of the cost model pictorially to provide a high-level overview and to ease exposition.

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 very 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. 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 $F_{kjt}(\cdot)$. These assumptions, and corresponding robustness analyses, are discussed at more length in the main text.

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 lognormal distribution, which is traditionally used to model medical expenditures, because we find that the lognormal distribution overpredicts 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:

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

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 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.

Table 13 presents summary results from the cost model estimation for the final choice model sample, including population statistics on the ACG index θ, the Weibull distribution parameters $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ for each category d, as well as the across category rank correlations $\rho_{z_{\theta_{it}}}(D, D')$. These are the fundamentals inputs used to generate F_{kjt}, as described above, and lead to very accurate characterizations of the overall total cost and out-of-pocket cost distributions (validation exercises which are not presented here).

Appendix B: Choice Model Estimation Algorithm Details

This appendix describes the details of the choice model estimation algorithm. The corresponding section in the text provided a high-level overview of this algorithm and outlined the estimation assumptions we make regarding choice model fundamentals and their links to observable data.

We estimate the choice model using a random coefficients probit simulated maximum likelihood approach similar to that summarized in Train (2009). The simulated maximum likelihood estimation approach has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. Since we use panel data, the likelihood function at the family level is computed for a *sequence* of choices from t_0 to t_2, 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.

First, the estimator simulates Q draws from the distribution of health expenditures output from the cost model, F_{kjt}, for each family, plan, and time period. These draws are used to compute plan expected utility conditional on all other preference parameters. It then simulates S draws for each family from the distributions of the random coefficients γ_k and δ_k, as well as from the distribution of the preference shocks $\epsilon_j(Y_k)$. We define the set of parameters θ as the full set of ex ante model parameters (before the S draws are taken):

$$\theta \equiv (\mu, \beta, \sigma^2_\gamma, \mu_\delta(Y_k), \sigma_\delta(Y_k), \alpha, \mu_{\epsilon_J}(Y_k), \sigma_{\epsilon_J}(Y_k), \eta_0, \eta_1, \eta_2).$$

We denote θ_{sk} one draw derived from these parameters for each family, including the parameters constant across draws:

$$\theta_{sk} \equiv (\gamma_k, \delta_k, \alpha, \epsilon_{JT}, \eta_0, \eta_1, \eta_2)$$

Denote θ_{Sk} the set of all S simulated draws for family k. For each θ_{sk} the estimator then uses all Q health draws to compute family-plan-time-specific expected utilities U_{skjt} following the choice model outlined in earlier in section 4. Given these expected utilities for each θ_{sk}, we simulate the probability of choosing plan j in each period using a smoothed accept-reject function with the form:

$$Pr_{skt}(j = j^*) = \frac{\left(\frac{\frac{1}{-U_{skj^*t}}(\cdot)}{\Sigma_J \frac{1}{-U_{skjt}}(\cdot)}\right)^\tau}{\Sigma_{\hat{j}}\left(\frac{\frac{1}{-U_{sk\hat{j}t}}(\cdot)}{\Sigma_J \frac{1}{-U_{skjt}}(\cdot)}\right)^\tau}$$

This smoothed accept-reject methodology follows that outlined in Train (2009) with some slight modifications to account for the expected utility specification. In theory, conditional on θ_{sk}, we would want to pick the j that maximizes U_{kjt} for each family, and then average over S to get final choice probabilities. However, doing this leads to a likelihood function with flat regions, because for small changes in the estimated parameters θ, the discrete choice made does not change. The smoothing function above mimics this process for CARA utility functions: as the smoothing parameter τ becomes large the smoothed Accept-Reject simulator becomes almost identical to the true Accept-Reject simulator just described, where the actual utility-maximizing option is chosen with probability one. By choosing τ to be large, an individual will always choose j^* when $\frac{1}{-U_{kj^*t}} > \frac{1}{-U_{kjt}} \forall j \neq j^*$. The smoothing function is modified from the logit smoothing function in Train (2009) for two reasons (i) CARA utilities are negative, so the choice should correspond to the utility with the lowest absolute value and (ii) the logit form requires exponentiating the expected utility, which in our case is already the sum of exponential functions (from CARA). This double exponentiating leads to computational issues that our specification overcomes, without any true content change since both models approach the true Accept-Reject function.

Denote any sequence of three choices made 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 s 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_{sk} a choice in period two will be made by a family k only if it is optimal conditional on θ_{sk}, other preference factors, and the switching costs implies by the period one choice. For all S simulation draws we compute the optimal sequence of choices for k with the smoothed Accept-Reject simulator, denoted j^3_{sk}. For any set of parameter values θ_{Sk} the probability that the model predicts j^3 will be chosen by k is:

$$\hat{P_k^{j^3}}(\theta, F_{kjt}, X_k^A, X_k^B, H_k, Y_k) = \Sigma_{s \in S} \mathbf{1}[j^3 = j^3_{sk}]$$

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

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

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 $SLL(\theta)$. In the results presented in the text, we choose $Q = 100$, $S = 50$, and $\tau = 4$, all values large enough such that the estimated parameters vary little in response to changes.

Appendix C: Moral Hazard Robustness analysis

In the text we discuss a robustness specification that investigates the cost model assumption of no moral hazard. To do this we necessarily make some substantial simplifying assumptions: for a full structural treatment of moral hazard in health insurance utilization see, e.g., Cardon and Hendel (2001), Einav et al. (2011) or Kowalski (2011). We implement the moral hazard robustness check by adjusting the output of the cost model to reflect lower total utilization in the less comprehensive plans (and vice-versa). The intent is to show that even when including price elasticities that are quite large relative to those found in the literature, the model output for switching costs and risk preferences does not change substantially. While the specification addresses moral hazard, it also sheds light on whether our estimates are sensitive to consumers having a reasonable level of private information above and beyond the detailed medical data we observe.

We operationalize this test in the following steps. First, we find the implied total spending changes across plans in the population for a price elasticity of -1.3, well higher than the range of -0.1 to -0.4 that represents most of the literature (for further discussion, see Chandra et al. (2010)). To do this we perform a back of the envelope calculation for the arc-elasticity of demand with respect to price, following the prior work of Einav et al. (2011) and Newhouse (1993). For this calculation, we use the average share of out-of-pocket spending for each plan as the price, and total medical expenditures as the quantity. For the three plans we study, the empirical shares of out-of-pocket spending are 15.5%, 20.9%, and 23.4% going from most to least comprehensive. We use these prices together with the average total spending of \$13,331 in PPO_{-1} at t_{-1} as the basis for this calculation. The formula for the arc-elasticity is:

$$Elasticity \ = \frac{(q_2 - q_1)/(q_2 + q_1)}{(p_2 - p_1)/(p_2 + p_1)}$$

We use the conservative elasticity of -1.3 and solve for the corresponding total cost changes this price response implies (here, p_j are the empirical shares of out-of-pocket spending for each plan and q_j is implied total spending in plan j). Solving for q_2 as a percentage of q_1 implies an, approximately, 25% reduction in total spending moving from PPO_{250} to PPO_{500}, a 33% reduction in total spending moving from PPO_{250} to $PP0_{1200}$, and a 10% reduction in total spending moving from PPO_{500} to PPO_{1200}. We then apply these reductions (or increases when moving into more comprehensive plans) and adjust the output of the cost model F_{kjt} according to the potential plan being chosen and the previous plan the family actually enrolled and incurred costs in. These new 'moral hazard adjusted' F_{kjt} then are input into the choice model, which is otherwise specified and estimated as in the text. The results from the analysis are presented in column 3 of table 7 and suggest that the initial assumption of no moral hazard does not markedly change our estimates of switching costs or risk preferences.[68]

It is important to point out that this exercise does not explicitly model the value of additional health spending, which the literature does through a non-linear budget constraint model where the family trades off the value of extra spending with the price of medical care (see Cardon and Hendel (2001), Einav et al. (2011) or Kowalski (2011)). Moral hazard here is captured purely by differences in spending. For a family that chooses PPO_{500} or PPO_{1200} in the data and is considering switching to PPO_{250}, our moral hazard cost wedge serves as an upper bound for the expected utility difference between these two plans and the more comprehensive plan, since their actual choice implies they value the increase in medical services less than the corresponding overall utility gain from a different financial lottery. The reverse is not true: for consumers considering switching to a less comprehensive plan our out-of-pocket cost wedge may not be an upper bound value of utility differences between two plans. However, we view this as a conservative approach in this case, since the high elasticity we've chosen together with the differences in the marginal prices of care between the plans (and the resulting implications for value of medical care foregone) make it unlikely that the value wedge between prospective plans is larger in reality than in this specification.

[68] In prior work, Handel (2010) presented evidence that the combined effect of moral hazard and selection on private information is not large in our setting. This suggests that the elasticity choice here of -1.3 represents more moral hazard than actually exists in our environment, implying this is a conservative approach.

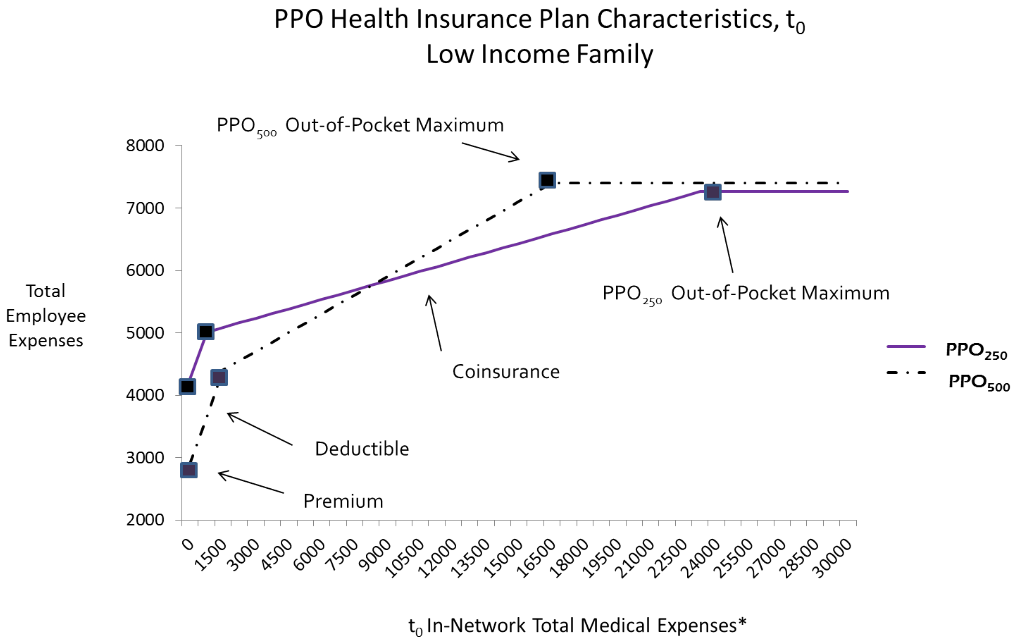

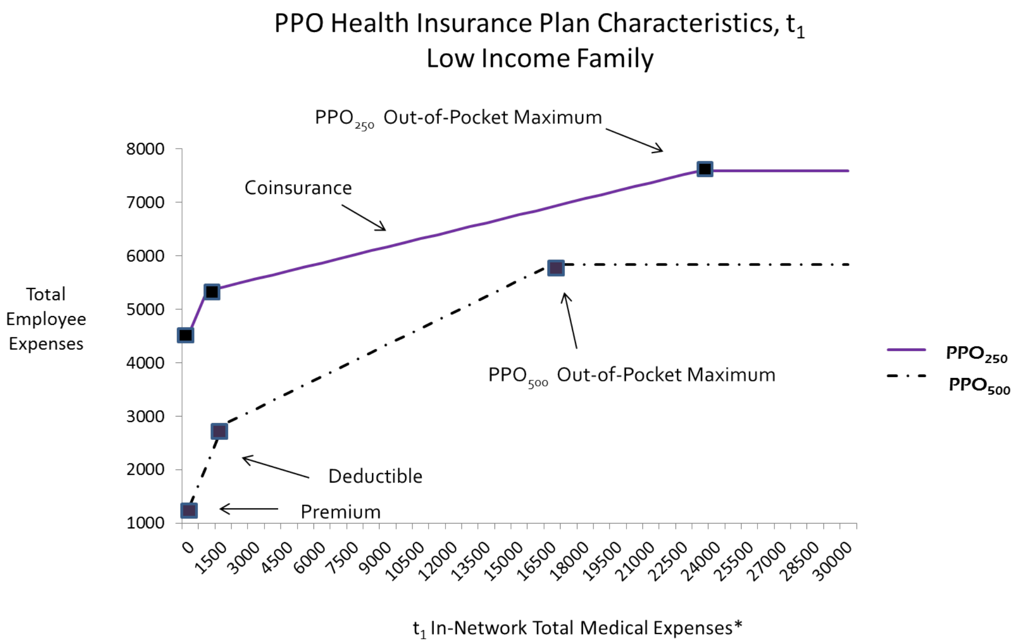

Figure 1: This figure describes the relationship between total medical expenses (plan plus employee) and employee out-of-pocket expenses in years t_0 and t_1 for PPO_{250} and PPO_{500}. This mapping depends on employee premiums, deductible, coinsurance, and out of pocket maximum. This chart applies to low income families (premiums vary by number of dependents covered and income tier, so there are similar charts for all 20 combinations of these two variables). Premiums are treated as pre-tax expenditures while medical expenses are treated as post-tax. 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 combination of income and dependents covered, 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.

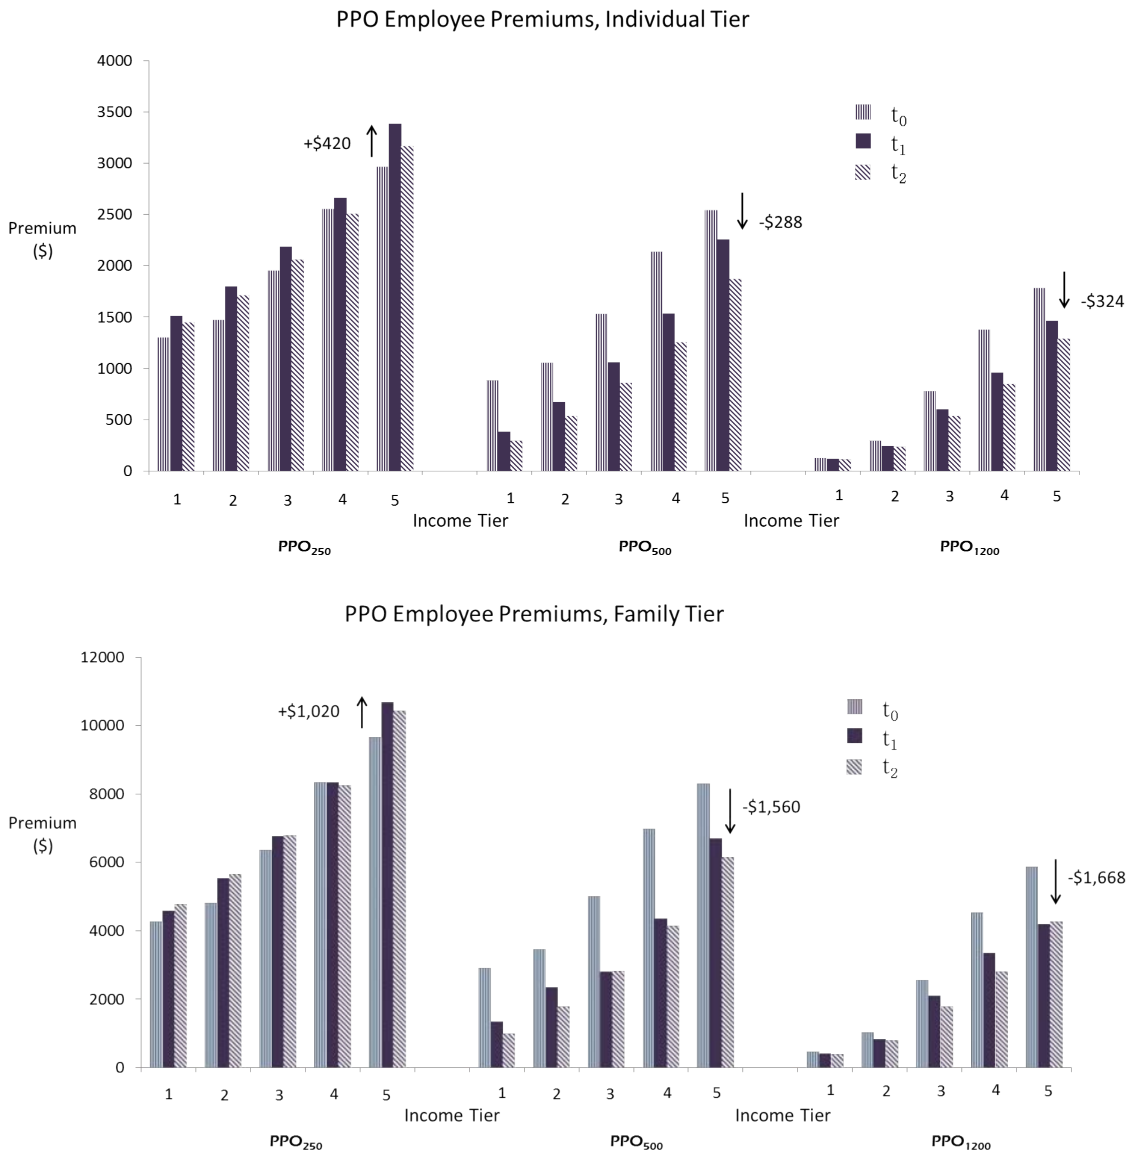

Figure 2: This figure describes the evolution of employee premium contributions at the firm over time between years t_0 and t_2. Employee premium contributions depend both on the number of dependents covered and the employee income tier. The top panel describes premiums for single employees and the bottom panel relates to families (employee + spouse + dependent(s)). The figure illustrates the large relative employee premium contribution changes between t_0 and t_1 across tiers. For example, a family in income tier five experienced a \$2,580 increase in PPO_{250} premium relative to PPO_{500}. The process for premium setting leading to this evolution over described is described and modeled in detail in section six.

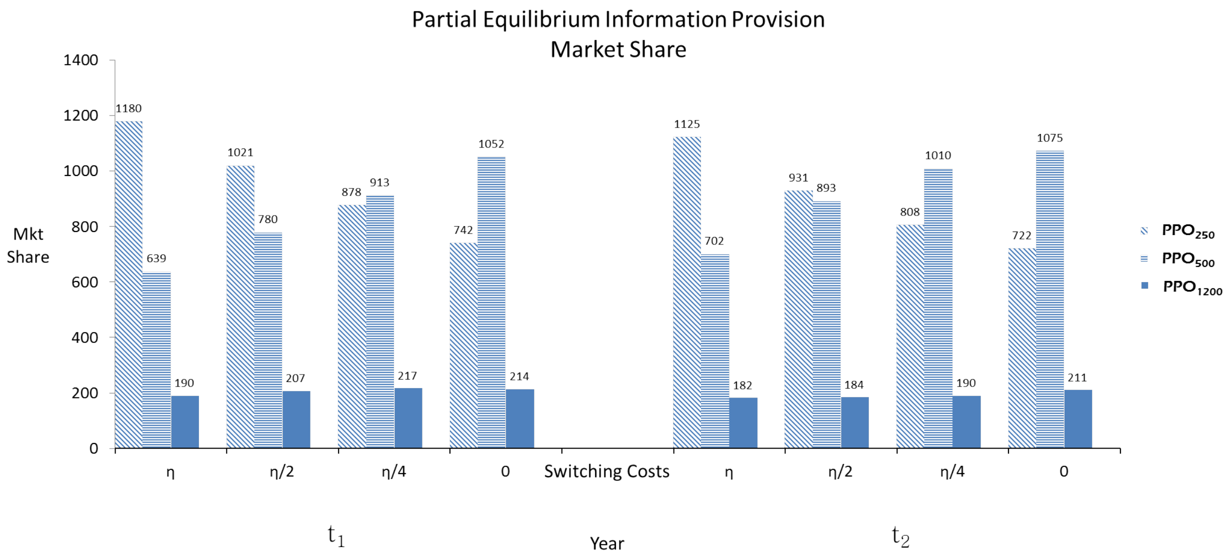

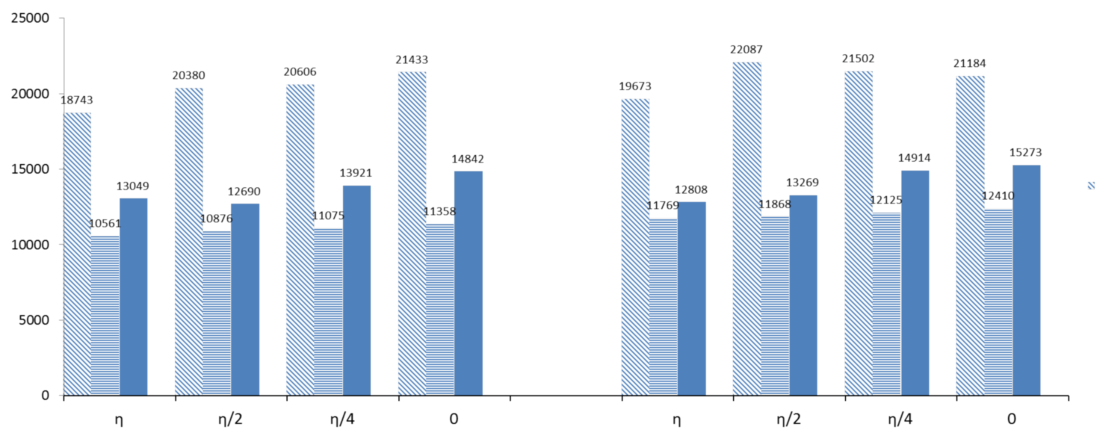

Figure 3: This figure describes the new plan enrollment levels and average costs under the 'naive' information provision counterfactual where consumer switching costs are reduced but prices are held fixed as observed in the data. The market shares are presented for all coverage tiers combined, while the average costs are presented for the family coverage tier (employee + spouse + dependent). The figure studies these quantities as a function of the effectiveness of the policy to reduce switching cost, where η signifies no reduction in switching costs from our estimates and 0 signifies a complete elimination of them. In general, the market share of PPO_{250} is decreasing in both t_1 and t_2 as the policy to reduce switching costs becomes more effective. The relative average cost of this plan increases as the market share declines, implying that when we allow for endogenous plan re-pricing there will be increased adverse selection on the margin.

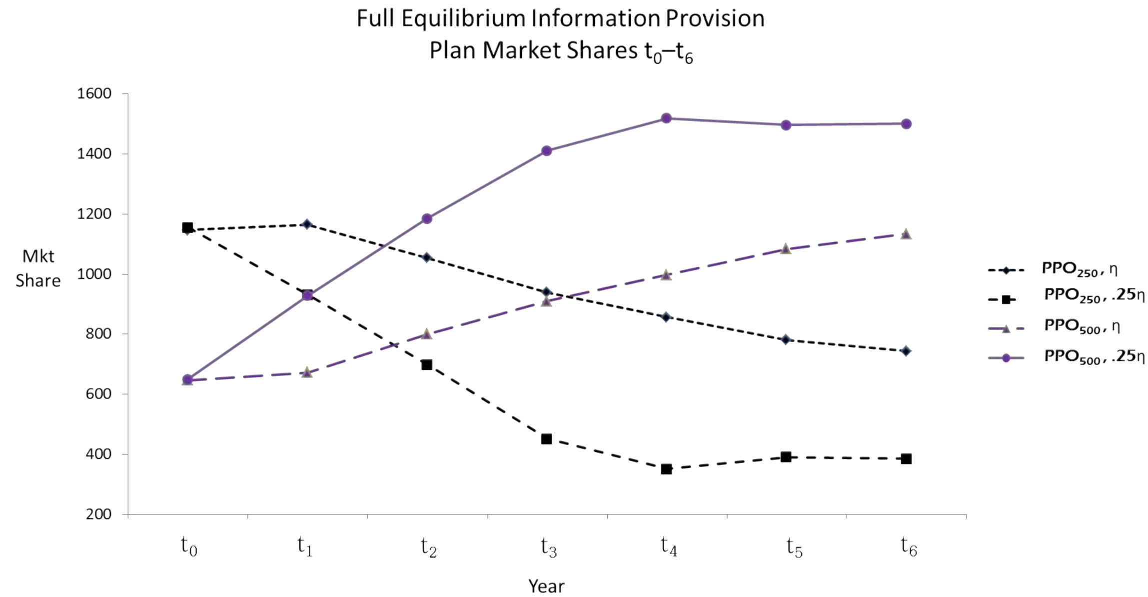

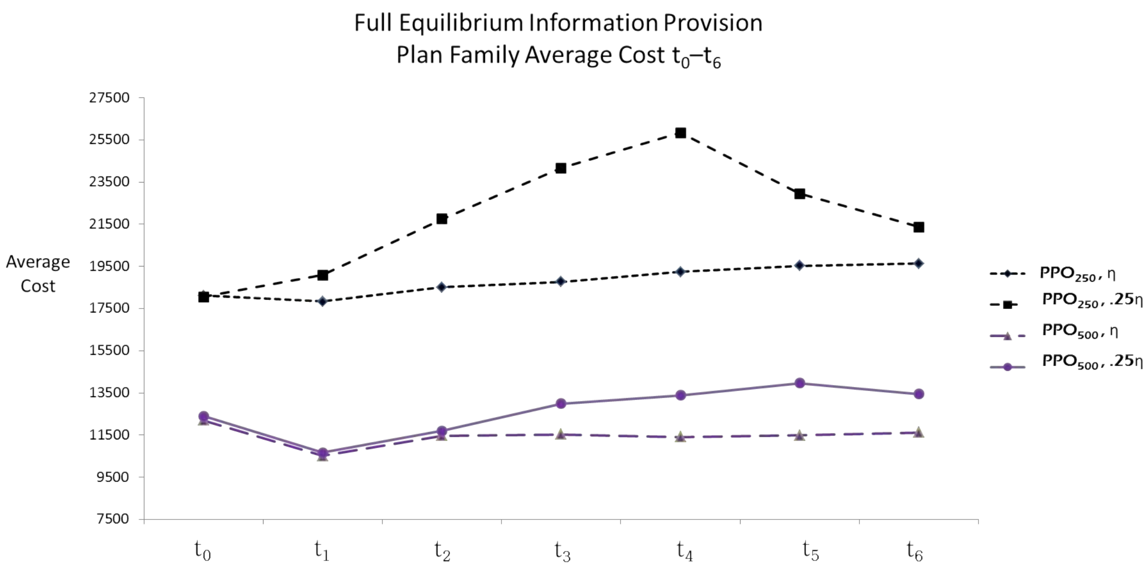

Figure 4: The top panel of this figure presents the time path of choices for PPO_{250} and PPO_{500} with and without the policy intervention to reduce switching costs. With endogenous plan pricing, the impact of the policy intervention on the market share of PPO_{250} relative to PPO_{500} is noticeable. In the benchmark case where there are significant switching costs η over the six year period the market share of PPO_{250} declines from 1147 to 744 while that of PPO_{500} increases from 647 to 1134. After the policy intervention reduces switching costs to $.25\eta$, PPO_{250} enrollment declines all the way to 385 after six years while PPO_{500} enrollment increases to 1501. In between t_0 and t_6, there are also noticeable differences in plan enrollment as a result of the policy intervention. The bottom panel in the figure shows the change in average costs for the family coverage tier under the policy intervention relative to the benchmark case of full switching costs. The average costs of PPO_{250} increase over time relative to those of PPO_{500}, signaling an increased relative premium for PPO_{250} and increased adverse selection.

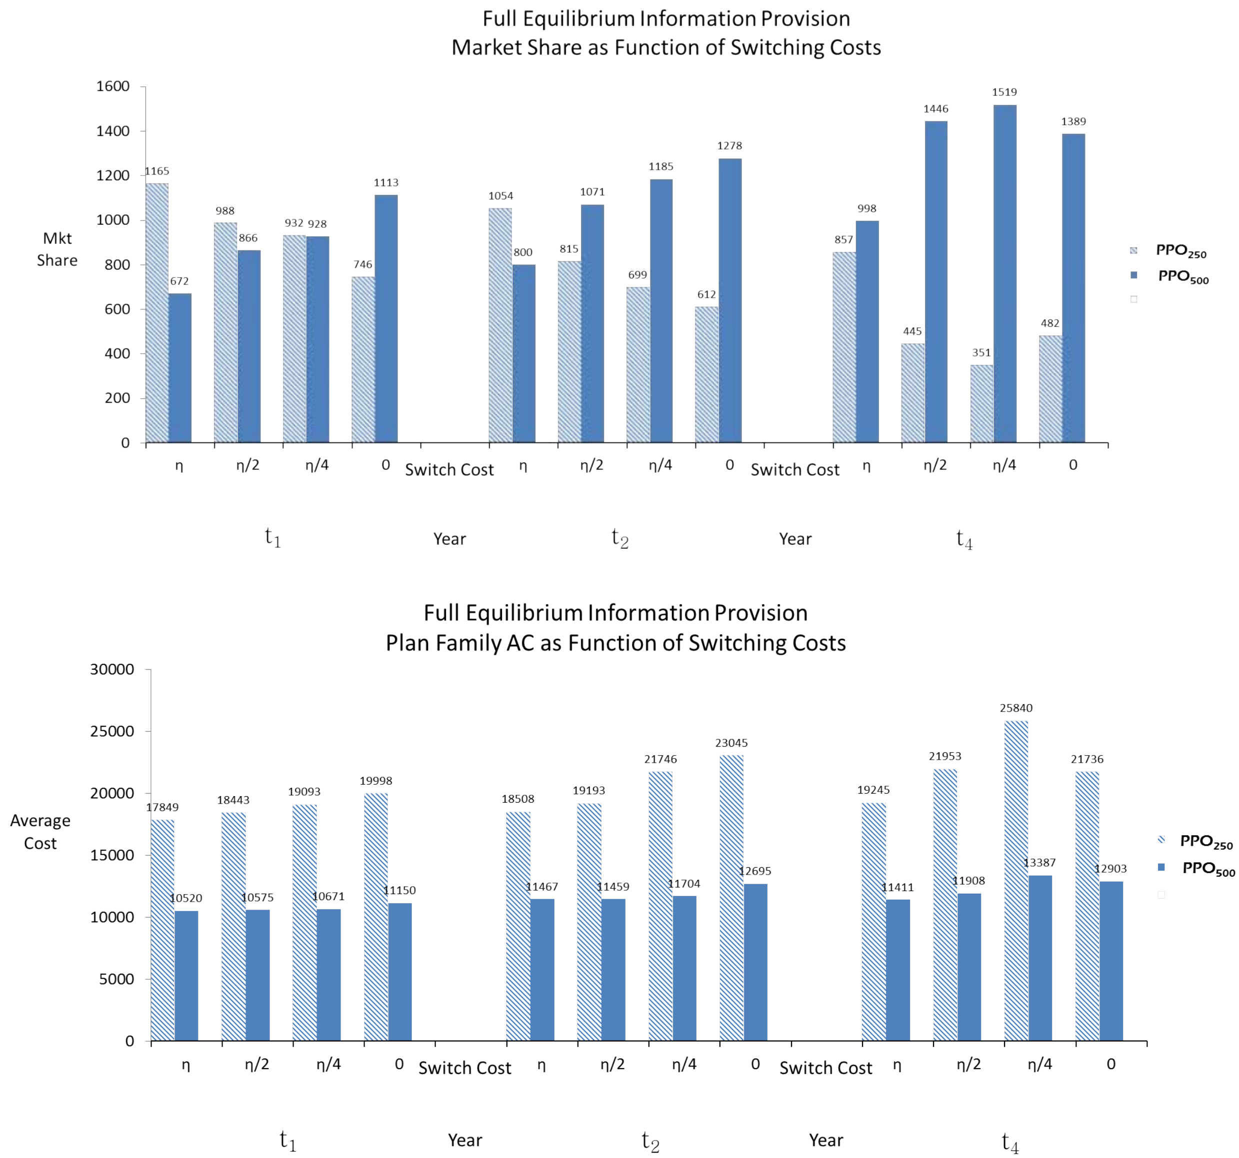

Figure 5: This figure shows market shares and average costs for $PP0_{500}$ and PPO_{250} in years t_1, t_2, and t_4, as a function of Z, in the environment with endogenous plan pricing. We report results for values of Z equal to 0,.25,.5, and 1. It is clear from the top panel that as Z decreases, and information provision becomes more effective, enrollment in PPO_{250} declines at the expense of enrollment in PPO_{500}. The bottom panel reveals that as Z decreases and enrollment increases, the average costs (and total premiums) of PPO_{250} increase relative to those of PPO_{500} in each year. This is indicative of increased relative premiums for comprehensive insurance and increased adverse selection.

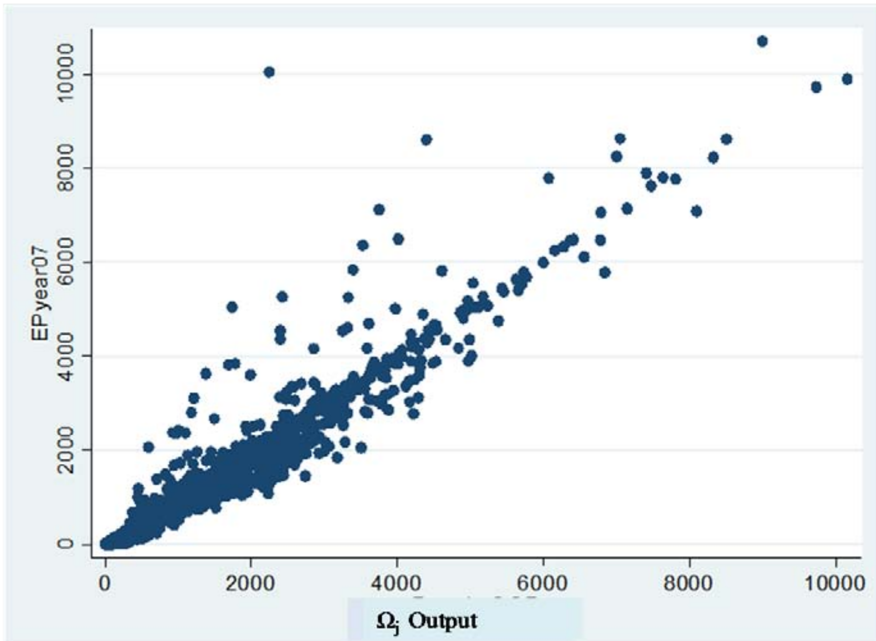

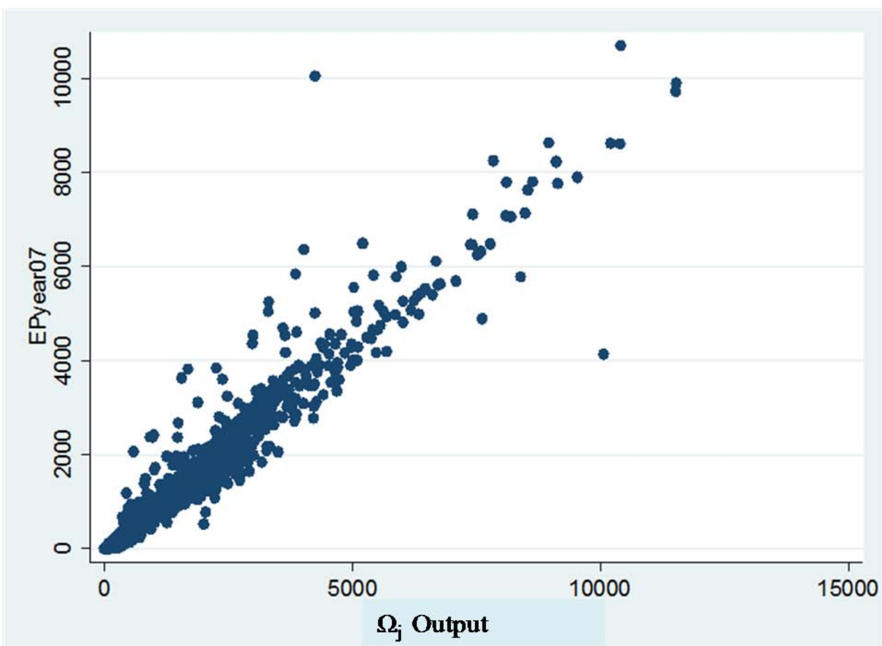

Figure 6: This figure validates the mapping Ω_j that translates the vector of total claims C into plan specific out of pocket expenditures. The two charts show Ω for PPO_{250}, with predicted out-of-pocket spending on the x-axis and actual out-of-pocket spending on the y-axis. The top chart is the mapping 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.

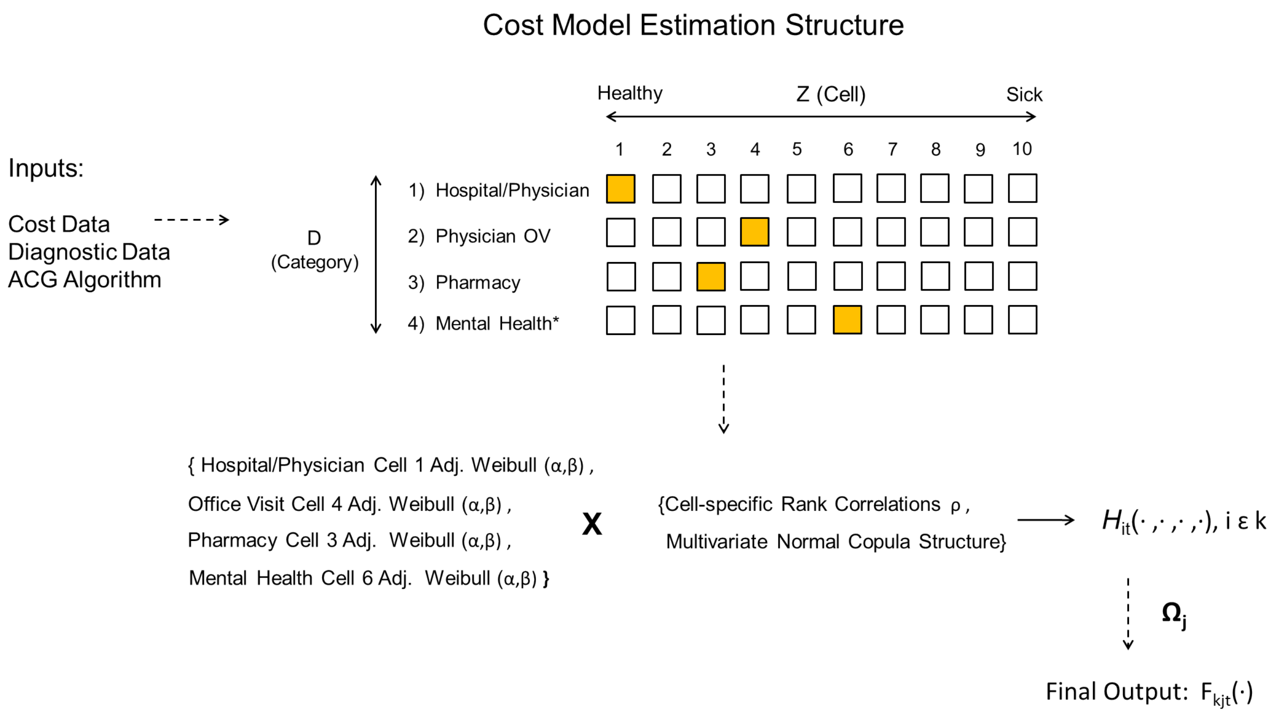

Figure 7: This figure outlines the primary steps of the cost model described in Appendix A.It moves from the initial inputs of cost data, diagnostic data, and the ACG algorithm to the final output F_{kjt} which is the family, plan, time specific distribution of out-of-pocket expenditures that enters the choice model for each family. The figure depicts an example individual in the top segment, corresponding to one cell in each category of medical expenditures. The last part of the model maps the expenditures for all individuals in one family into the final distribution F_{kjt}.

Sample Demographics	All Employees	PPO Ever	Final Sample
N - Employee Only	11,253	5,667	2,023
N - All Family Members	20,963	10,713	4,544
Mean Employee Age (Median)	40.1 (37)	40.0 (37)	42.3 (44)
Gender (Male %)	46.7%	46.3%	46.7%
Income			
Tier 1 (< $41K)	33.9%	31.9%	19.0%
Tier 2 ($41K-$72K)	39.5%	39.7%	40.5%
Tier 3 ($72K-$124K)	17.9%	18.6%	25.0%
Tier 4 ($124K-$176K)	5.2%	5.4%	7.8%
Tier 5 (> $176K)	3.5%	4.4%	7.7%
Family Size			
1	58.0 %	56.1 %	41.3 %
2	16.9 %	18.8 %	22.3 %
3	11.0 %	11.0 %	14.1 %
4+	14.1 %	14.1 %	22.3 %
Staff Grouping			
Manager (%)	23.2%	25.1%	37.5%
White-Collar (%)	47.9%	47.5%	41.3%
Blue-Collar (%)	28.9%	27.3%	21.1%
Additional Demographics			
Quantitative Manager	12.8%	13.3%	20.7%
Job Tenure Mean Years (Median)	7.2 (4)	7.1 (3)	10.1 (6)
Zip Code Population Mean (Median)	42,925 (42,005)	43,319 (42,005)	41,040 (40,175)
Zip Code Income Mean (Median)	$56,070 ($55,659)	$56,322 ($55,659)	$60,948 ($57,393)
Zip Code House Value Mean (Median)	$226,886 ($204,500)	$230,083 ($209,400)	$245,380 ($213,300)

Table 1: This table presents summary demographic statistics for the population we study. The first column describes demographics for the entire sample whether or not they ever enroll in insurance with the firm. The second column summarizes these variables for the sample of individuals who ever enroll in a PPO option, the choices we focus on in the empirical analysis. 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_1. Comparing the columns shows little selection on demographics into PPO options and some selection based on family size into the final estimation sample.

PPO Plan Characteristics	PPO_{250}	PPO_{500}	PPO_{1200}
Basic Characteristics			
Individual Deductible (Family)	250 (750)	500 (1500)	1200 (2400)
Coinsurance Rate	10%	20%	20%
Physician Office Visit Copay	25	25	NA
Emergency Room Copay	100	100	NA
Mental Health Coinsurance	50%	50%	50%
Prescription Drug Copay	5/25/45*	5/25/45*	NA
	(10/50/75)	(10/50/75)	NA
Individual OOP Max (Family)**			
Income Tier 1 (< \$41K)	1000 (3000)	1500 (4500)	2000 (6000)
Income Tier 2-3 (\$41K-\$124K)	2000 (5000)	3000 (7000)	4000 (8000)
Income tier 4-5 (> \$124K)	3000 (8000)	4000 (9000)	5000 (10000)
Other Features			
Free Preventative Care	No	No	Yes
Traditional FSA Eligible	Yes	Yes	No
HSA Eligible	No	No	Yes

* Prescription copay max of 1500 per person.
Copays for 30-day supply. 90-day supply copay in parentheses.
Copays increase in t_2 by approx. 20%
** 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 available at the firm after the menu change at t_0. 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 up to 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 the out-of-pocket maximum for each plan. Out-of-pocket maximums vary with income, presumably for equity considerations. This chart does not include out-of-network plan characteristics which account for only 2% of total expenses and don't vary substantially across the three plans. These three plans have the same provider network and cover the same medical services.

Choice Behavior			
	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%)

Switching Behavior		
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	22
HMO_2	9	6	5	6	983

Table 3: This table describes plan choice at the firm between years t_{-1} and t_1. The top panel documents the number of employees who chose each plan in each of these three years. Before the menu change about 39% of employees enrolled in a PPO and 47% in an HMO. After the menu change these numbers are essentially reversed with 47% in a PPO and 35% in an HMO. The middle panel studies the choice behavior of all employees at the firm who were enrolled in any plan in both of the years t_{-1} and t_0. Each column corresponds to the plan an employee was in at t_{-1} while each row corresponds to the plan an employee was in at 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. These panels together reveal significant horizontal differentiation between the nest of PPOs and nest of HMOs.

New Enrollee Analysis	New Enrollee t_{-1}	New Enrollee t_0	New Enrollee t_1
N, t_0	1056	1377	-
N, t_1	784	1267	1305
t_0 **Choices**			
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 **Choices**			
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
Female %	56%	54%	53%
Manager %	20%	18%	19%
FSA Enroll %	15%	12%	14%
Dental Enroll %	88%	86%	86%
Median (Mean) Expense t_1	844 (4758)	899 (5723)	-
Income Tier 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 4: This table describes the choice behavior of new employees at the firm over several consecutive years and presents our first model-free test of switching costs. Each column describes one cohort of new employees at the firm, corresponding to a specific year of arrival. First, the chart describes the health insurance choices made by these cohorts in year t_0 (the year of the insurance plan menu change) and in the following year, t_1. The last part of the chart lists the demographics for each cohort of new arrivals at the time of their arrival. Given the very similar demographic profiles and large sample size for each cohort, if there are no switching costs, the t_1 choices of employees who entered the firm at t_0 and t_{-1} should be identical to the t_1 choices of employees who entered the firm at t_1. The table shows that, in fact, the active choices made by the t_1 cohort are quite different than those of the prior cohorts in the manner we would expect with high switching costs: the t_1 choices of employees who enter at t_0 and t_{-1} reflect both t_1 prices and t_0 choices while the t_1 choices of new employees at t_1 reflect 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 illustrates the large impact that switching costs have on choices in our setting, independent of the choice model setup and structure.

Dominated Plan Analysis	t_1 Dominated Stay	t_1 Dominated Switch	t_2 Dominated Stay	t_2 Dominated Switch
N	498	61	378	126
Minimum Money Lost*	$374	$453	$396	$306
PPO_{500}	-	44 (72%)	-	103 (81%)
PPO_{1200}	-	4 (7%)	-	6(5%)
Any HMO	-	13 (21%)	-	17 (14%)
FSA t_1	25.4%	32.1%	27.2%	28.6%
FSA t_2	-	-	28.1%	30.9%
Dental Switch t_1	4.3%	14.1%	3.5%	10.9%
Dental Switch t_2	-	-	6.9%	17.2%
Age (mean)	44.9	38.3	46.2	41.4
Income Tier (mean)**	1.6	1.4	1.6	1.7
Quant. Manager	11%	8%	11%	11%
Single (%)	40%	41%	40%	33%
Male (%)	42%	46%	39%	55%

All Plan Analysis	PPO_{250} Stay t_1	PPO_{250} Switch t_1	All Plans t_1 Stay	All Plans t_1 Switch
Sample Size	1626	174	2786	384
FSA 2008 Enrollee	31%	41%	25%	39%
Dental Switch	3.2%	13.1%	3.8%	14.5%
Age (mean)	48.3	40.6	44.0	39.1
Income Tier (mean)**	2.5	2.2	2.3	2.1
Quant. Manager	20%	17%	17%	14%
Single (%)	50%	56%	53%	59%
Male (%)	48%	42%	49%	40%

Table 5: This top panel in 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 (498 out of 559 (89%)) remain in PPO_{250} even after it becomes dominated by PPO_{500} with 378 of 504 (25%) still remaining in this plan at t_2. People who do switch are more likely to exhibit a pattern of active choice behavior in general as evidenced 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. The bottom panel studies the profiles of those who switch at t_1 and those who don't for the two groups of (i) PPO_{250} enrollees at t_0 and (ii) the entire universe of PPO plan enrollees present in t_0 and t_1. This reveals a similar pattern of active decision making as switchers in these populations are also more likely to enroll in FSAs and switch dental plans.

Final Sample Total Expenses	PPO_{-1}	PPO_{250}	PPO_{500}	PPO_{1200}
Family t_{-1} Total Expenses ($)				
t_{-1}				
N Employees (Mean Family Size)	2,022 (2.24)	-	-	-
Mean (Median)	13,331 (4916)	-	-	-
25th pct.	1,257	-	-	-
75th pct.	13,022	-	-	-
t_0				
N (Mean Family Size)	-	1328 (2.18)	414 (2.20)	280 (2.53)
Mean (Median)	-	16,976 (6,628)	6,151 (2,244)	6,742 (2,958)
25th pct.	-	2,041	554	658
75th pct.	-	16,135	6,989	8,073
t_1				
N (Mean Family Size)	-	1,244 (2.19)	546 (2.19)	232 (2.57)
Mean (Median)	-	17,270 (6,651)	7,759 (2,659)	6,008 (2,815)
25th pct.	-	2,041	708	589
75th pct.	-	16,707	8,588	7,191
Individual Category Expenses				
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 6: This table investigates the extent of adverse selection across PPO options after the t_0 menu change for those in the final estimation sample. 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. The numbers in the table for all choices represent t_{-1} total claims in dollars so that these costs can proxy for health risk without being confounded by moral hazard (t_0 and t_1 costs could be the result of selection or moral hazard). The table reveals that those who choose PPO_{250} have much higher expenditures at t_{-1} than those who choose the other two plans, implying substantial selection on observables in the vein of Finkelstein and Poterba (2006). The bottom panel presents a breakdown of these costs according to our cost model expenditure categories.

Empirical Model Results

Parameter	Primary	Base	MH Robust	γ Robust	ϵ Robust
Switching Cost - Single, η_0	1729	1779	1859	2430	1944
	(28)	(72)	(107)	(116)	(150)
Switching Cost - Family, $\eta_0 + \eta_2$	2480	2354	2355	3006	2365
	(26)	(62)	(113)	(94)	(34)
SC - FSA Enroll, η_1	-551	-	-669	-723	-417
	(56)	-	(155)	(131)	(50)
SC - Income, η_1	-32	-	-59	-8	-7
	(13)	-	(15)	(43)	(15)
SC - Quantitative, η_1	5	-	-40	-537	-6
	(138)	-	(80)	(223)	(92)
SC - Manager, η_1	198	-	277	875	224
	(292)	-	(164)	(200)	(244)
SC - Chronic Condition, η_1	80	-	29	-221	67
	(46)	-	(67)	(148)	(35)
SC - Salient Change, η_1	156	-	95	61	123
	(83)	-	(60)	(212)	(54)
SC - PPO_{1200}, η_1	-19	-	-32	-327	-113
	(184)	-	(46)	(122)	(52)
SC - Total Pop. Mean, η	2032	2087	1886	1914	1986
[Pop. Standard Deviation]	[446]	[286]	[387]	[731]	[316]
Risk Aversion Mean - Intercept , μ_γ	$2.32 * 10^{-4}$	$3.12 * 10^{-4}$	$2.31 * 10^{-4}$	-8.94	$1.90 * 10^{-4}$
	$(9.0 * 10^{-6})$	$(1.1 * 10^{-5})$	$(1.1 * 10^{-5})$	(0.43)	$(1.0 * 10^{-5})$
Risk Aversion Mean - Income ,β	$2.90 * 10^{-5}$	$4.21 * 10^{-5}$	$1.80 * 10^{-5}$	0.07	$2.40 * 10^{-5}$
	$(4.0 * 10^{-6})$	$(3.0 * 10^{-6})$	$(3.0 * 10^{-6})$	(0.016)	$(3.0 * 10^{-6})$
Risk Aversion Mean - Age , β	$2.27 * 10^{-6}$	-	$3.45 * 10^{-6}$	0.28	$2.59 * 10^{-6}$
	$(1.7 * 10^{-7})$	-	$(1.8 * 10^{-7})$	(0.011)	$(1.5 * 10^{-7})$
Risk Aversion Std. Deviation , σ_γ	$1.88 * 10^{-4}$	$1.19 * 10^{-4}$	$1.27 * 10^{-4}$	1.37	$1.04 * 10^{-4}$
	$(6.6 * 10^{-5})$	$(8.0 * 10^{-6})$	$(6.0 * 10^{-6})$	(0.06)	$(5.9 * 10^{-5})$
CDHP - Single - RC Mean , δ	-2912	-3665	-2801	-2985	-2833
	(754)	(218)	(416)	(85)	(130)
CDHP - Single - RC Variance, σ_δ	843	1283	1070	989	1141
	(431)	(104)	(139)	(70)	(113)
CDHP - Family - RC Mean , δ	-2871	-4847	-2614	-5344	-2932
	(73)	(204)	(115)	(134)	(40)
CDHP - Family - RC Variance, σ_δ	897	1733	1149	2179	1013
	(28)	(99)	(132)	(80)	(31)
High Total Cost - PPO_{250}, α	856	937	607	1386	860
	(50)	(175)	(55)	(264)	(66)
ϵ_{500} - Single	204	60	51	50	-
	(13)	(67)	(30)	(55)	-
ϵ_{1200} - Single	502	872	647	161	-
	(475)	(122)	(228)	(72)	-
ϵ_{500} - Family	329	180	789	90	-
	(25)	(94)	(28)	(89)	-
ϵ_{1200} - Family	811	888	715	676	-
	(25)	(104)	(44)	(426)	-

Table 7: This table presents the estimated parameter results for the primary choice model from section 4 and the four robustness checks outlined in section 5. All non-risk aversion coefficients are in dollar units with standard errors for parameters given in parentheses. The results from the Primary specification are the inputs into the counterfactual simulations presented in section six.

Risk Preference Analysis	Absolute Risk Aversion	Interpretation
Normal Heterogeneity		
Mean / Median Individual	$4.22 * 10^{-4}$	94.6
25th percentile	$2.95 * 10^{-4}$	96.1
75th percentile	$5.49 * 10^{-4}$	93.8
95th percentile	$7.31 * 10^{-4}$	92.2
99th percentile	$8.59 * 10^{-4}$	91.8
Log normal Heterogeneity		
Mean	$9.82 * 10^{-4}$	91.0
25th percentile	$1.53 * 10^{-4}$	97.2
Median	$3.85 * 10^{-4}$	95.0
75th percentile	$9.72 * 10^{-4}$	91.1
95th percentile	$3.70 * 10^{-3}$	72.8
99th percentile	$9.30 * 10^{-3}$	51.1
Comparable Estimates		
Cohen and Einav (2007) Benchmark Mean	$3.1 * 10^{-3}$	76.5
Cohen and Einav (2007) Benchmark Median	$3.4 * 10^{-5}$	99.7
Gertner (1993)	$3.1 * 10^{-4}$	97.0
Holt and Laury (2002)	$3.2 * 10^{-2}$	21.0
Sydnor (2010)	$2.0 * 10^{-3}$	83.3

Table 8: This table examines the risk preference estimates from the empirical results presented in table 7. The first section of the table is for the normally distributed risk preference estimates in the Primary specification, where the age and income coefficients are evaluated at the median values of those variables. The second section is for the model with log-normally distributed preferences studied in column 4 of table 7. 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 (wide) range found in the literature and show moderate risk aversion except at the tails in the log-normal model where consumers are quite risk averse.

Partial Equilibrium Welfare Analysis		
	t_1	t_2
Mean Δ CEQ		
Population	\$96	\$114
Switchers Only	\$175	\$196
Mean Welfare Change: % Total Premiums		
Mean Employee Premium (MEP)	\$2,067	\$1,954
Welfare Change Population	4.6%	5.8%
Welfare Change Switchers	8.5%	10.0%
Mean Welfare Change: % Total Emp. Spending		
Mean Total Emp. Spending	\$4,373	\$4,486
Welfare Change Population	2.2%	2.5%
Welfare Change Switchers	4.0%	4.4%
Mean Welfare Change: % $\|CEQ\|$		
Mean Total $\|CEQ\|$	\$6,694	\$6,773
Welfare Change Population	1.4%	1.7%
Welfare Change Switchers	2.6%	2.9%

Table 9: This table presents the welfare results of the 'naive' counterfactual information provision analysis. We present the mean per employee per year dollar change in certainty equivalents and corresponding percentage welfare changes resulting from the policy intervention that reduces switching costs to $.25\eta$ from η. We present three alternative welfare benchmarks to assess the % impact of this these certainty equivalent changes. These metrics divide the change in certainty equivalent from the policy intervention by mean (i) total employee premiums (ii) total employee spending and (iii) the absolute value of the certainty equivalent loss. Note that since all figures are losses the certainty equivalent absolute value is larger than the total spending figure. Since we hold plan prices fixed in this exercise, the welfare changes must be positive when consumers make better individual-level decisions.

Full Equilibrium Welfare Analysis **Information Provision: η to .25η**	t_1	t_2	t_4	t_6	**Avg.** t_1-t_6
Mean Δ CEQ					
Population	-$63	-$104	-$144	-$118	-$115
Switcher Pop. %	51%	49%	48%	53%	49%
Switchers Only	$86	$175	$ 245	$242	$186
Non-Switchers Only	-$205	-$391	-$555	-$432	-$442
High Expense Pop.[a] %	10%	11%	11%	11%	11%
High Expense	$26	$106	$119	$65	$62
Non-High Expense	-$73	-$130	-$177	-$141	-$137
Single Pop. %	47%	46%	46%	46%	46%
Single	-$249	-$367	-$414	-$195	-$319
w/ Dependents	$99	$124	$89	-$51	$61
Low Income Pop.[b] %	40%	41%	41%	41%	41%
Low Income	-$81	-$218	-$282	-$178	-$200
High Income	-$36	$62	$57	-$30	$0
Welfare Change: % Premiums[c]					
Mean Employee Premium	$1,471	$1,591	$1,455	$1,259	$1,500
Welfare Change Population	-4.8%	-6.5%	-9.9%	-9.4%	-7.7%
Welfare Change Switchers	5.6%	11.0%	16.9%	19.2%	12.4%
Welfare Change Non-Switchers	-13.9%	-24.6%	-38.1%	-34.3%	-29.4%
Welfare Change: % Total Spending[d]					
Mean Total Emp. Spending	$3,755	$4,097	$4,022	$3,862	$4,015
Welfare Change Population	-1.7%	-2.5%	-3.6%	-3.06%	-2.9%
Welfare Change Switchers	2.3%	4.3%	6.1%	6.3%	4.6%
Welfare Change Non-Switchers	-5.5%	-9.5%	-13.8%	-11.2%	-11.0%
Welfare Change: % $\|CEQ\|$[e] Loss					
Mean Total $\|CEQ\|$	$5,888	$6,264	$6,207	$6,065	$6,190
Welfare Change Population	-1.1%	-1.7%	-2.3%	-2.0%	-1.9%
Welfare Change Switchers	1.5%	2.8%	4.0%	4.0%	3.0%
Welfare Change Non-Switchers	-3.5%	-6.2%	-8.9%	-7.1%	-7.1%

Table 10: This table presents the welfare results of the endogenous insurance pricing policy counterfactual for the case where switching costs are reduced from η to 0.25η. We present the change in the mean per employee per year certainty equivalent moving from the simulation with full switching costs to reduced switching costs. In addition to studying the effect of the policy on efficiency, we study the distributional effects based on four 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 (iii) whether an employee is single or covers dependents and (iv) whether an employee has high or low income. We present the same welfare metrics as in table 9, 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 loss.

Full Equilibrium Welfare Analysis **Information Provision Range**	First-Best	Baseline	.75η	.5η	.25η	0
Mean Δ CEQ **(% of Premiums)**						
Population	\$123 (8.2%)	- (-)	-\$41 (-2.7%)	-\$73 (-4.9%)	-\$115 (-7.7%)	-\$107 (-7.1%)
Switchers	-\$538 (-35.9%)	- (-)	\$1,017 (67.8%)	\$766 (51.0%)	\$186 (12.4%)	\$118 (7.9%)
Non-Switchers	\$953 (63.5%)	- (-)	-\$249 (-16.6%)	-\$371 (-24.8%)	-\$442 (-29.4%)	-\$382 (-25.4%)
High Expense	\$936 (62.4%)	- (-)	\$38 (2.6%)	\$84 (5.6%)	\$62 (4.2%)	\$121 (8.1%)
Non-High Expense	\$22 (1.5%)	- (-)	-\$52 (-3.5%)	-\$93 (-6.2%)	-\$137 (-9.2%)	-\$136 (-9.1%)
Single	-\$683 (-45.5%)	- (-)	-\$153 (-10.2%)	-\$295 (-19.7%)	-\$319 (-21.2%)	-\$286 (-19.0%)
Family	\$826 (55%)	- (-)	-\$54 (3.6%)	\$119 (7.9%)	\$61 (4.1%)	\$47 (3.1%)
Low Income	-\$349 (-23.3%)	- (-)	-\$75 (-5.0%)	-\$153 (-10.2%)	-\$200 (-13.3%)	-\$190 (-12.7%)
High Income	\$806 (53.7%)	- (-)	\$10 (0.6%)	\$43 (2.9%)	\$0 (0)	\$13 (0.9%)

Table 11: This table shows the welfare change of a range of policy interventions, in terms of effectiveness, relative to the baseline where preferences are as estimated in table 7. In addition, we present results on the welfare loss from adverse selection in the actual environment relative to the first-best. The chart reports the change in the mean per employee per year certainty equivalent in each environment, relative to the baseline case. In parentheses, we include the percentage corresponding to this certainty equivalent change divided by mean employee premiums paid per employee per year. Column 1 shows how the first-best compares to the baseline and reveals that the mean welfare loss from adverse selection in the current information environment is \$123 or 8.2% of total premiums paid in the baseline. Columns 3 through 6 correspond to different counterfactual environments where switching costs have been reduced relative to the baseline. We study four cases, when switching costs are assumed to be 75%, 50%, 25%, and 0% of baseline switching costs respectively. We report welfare results for the population as well as different segments of the population. The 25% counterfactual is examined in more detail in table 10.

Full Equilibrium Information Provision Switching Costs Analysis		η	$.75\eta$	$.5\eta$	$.25\eta$	0
t_1	SC / Switcher	1,982	1,502	1,010	506	0
	Switcher %	14%	19%	21%	23%	32%
	Avg. SC Pop.	291	284	216	118	0
t_2	SC / Switcher	1,959	1,488	1001	506	0
	Switcher %	14%	17%	20%	23%	25%
	Avg. SC Pop.	269	252	201	117	0
t_4	SC / Switcher	1,973	1,506	982	480	0
	Switcher %	7%	10%	11%	12%	15%
	Avg. SC Pop.	144	156	109	56	0
t_6	SC / Switcher	1,942	1,451	962	483	0
	Switcher %	6%	8%	10%	12%	16%
	Avg. SC Pop.	110	109	90	60	0
Avg. t_1-t_6	SC / Switcher	1,963	1,489	988	493	0
	Switcher %	9%	13%	14%	17%	20%
	Avg. SC Pop.	185	188	142	83	0

Welfare Impact		η	$.75\eta$	$.5\eta$	$.25\eta$	0
$\kappa = 0$	Welfare Relevant SC	0	0	0	0	0
	Δ CEQ (% Premiums)	-	-$41 (-2.7%)	-$73 (-4.9%)	-$115 (-7.7%)	-$107 (-7.1%)
$\kappa = 0.25$	Welfare Relevant SC	46	47	36	21	0
	Δ CEQ (% Premiums)	-	-$42 (-2.8%)	-$63 (-4.2%)	-$90 (-6.0%)	-$61 (-4.1%)
$\kappa = 0.5$	Welfare Relevant SC	93	94	71	42	0
	Δ CEQ (% Premiums)	-	-$42 (-2.8%)	-$51 (-3.4%)	-$64 (-4.3%)	-$14 (-0.9%)
$\kappa = 1$	Welfare Relevant SC	185	188	142	83	0
	Δ CEQ (% Premiums)	-	-$44 (-2.9%)	-$30 (-2.0%)	-$13 (-0.9%)	-$78 (5.2%)

Table 12: Table 12 expands the welfare analysis to account for the possibility that some proportion of estimated, and subsequently reduced, switching costs should be included in the welfare analysis. Tables 10 and 11 present results conditional on $\kappa = 0$ (switching costs are not welfare relevant) while this table presents results across the range of κ from 0 to 1 (all switching costs are welfare relevant). The top panel of this table studies the profile of switching costs incurred for different Z from t_1 to t_6, while the bottom panel assesses the welfare impact of these interventions as a function of κ. The table reveals that for almost all combinations of κ and Z, there is a negative welfare impact from reduced switching costs and better consumer decisions.

Final Sample **Cost Model Output**	Overall	PPO_{250}	PPO_{500}	PPO_{1200}
Individual Mean (Median)				
Unscaled ACG Predictor				
Mean		1.42	0.74	0.72
Median		0.83	0.37	0.37
Pharmacy: Model Output				
Zero Claim Pr.	0.35 (0.37)	0.31 (0.18)	0.40 (0.37)	0.42 (0.37)
Weibull α	1182 (307)	1490 (462)	718 (307)	596 (307)
Weibull β	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)
Mental Health				
Zero Claim Pr.	0.88 (0.96)	0.87 (0.96)	0.90 (0.96)	0.90 (0.96)
Weibull α	1422 (1295)	1447 (1295)	1374 (1295)	1398 (1295)
Weibull β	0.98 (0.97)	0.99 (0.97)	0.98 (0.97)	0.98 (0.97)
Hospital / Physician				
Zero Claim Pr.	0.23 (0.23)	0.21 (0.23)	0.26 (0.23)	0.26 (0.23)
Weibull α	2214 (1599)	2523 (1599)	1717 (1599)	1652 (1599)
Weibull β	0.58 (0.55)	0.59 (0.55)	0.55 (0.55)	0.55 (0.55)
($> \$40,000$) Claim Pr.	0.02 (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)
Physician OV				
Zero Claim Pr.	0.29 (0.20)	0.26 (0.20)	0.33 (0.46)	0.34 (0.46)
Weibull α	605 (553)	653 (553)	517 (410)	529 (410)
Weibull β	1.15 (1.14)	1.15 (1.14)	1.15 (1.14)	1.14 (1.14)
Correlations				
Rank Correlation Hospital-Pharm.	0.28 (0.34)	0.26 (0.32)	0.31 (0.34)	0.32 (0.34)
Rank Correlation Hospital-OV	0.73 (0.74)	0.72 (0.74)	0.74 (0.74)	0.74 (0.74)
Rank Correlation Pharm.-OV	0.35 (0.41)	0.33 (0.37)	0.38 (0.41)	0.39 (0.41)

Table 13: This table describes the output of the cost model in terms of the means and medians of individual level parameters, classified by the plan actually chosen. These parameters are aggregated for these groups but have more micro-level groupings, which are the primary inputs into our cost projections in the choice model. Weibull α, Weibull β, and Zero Claim Probability correspond to the cell-specific predicted total individual-level health expenses as described in more detail in Appendix A.