Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare *

Benjamin R. Handel
Economics Department, UC Berkeley and NBER

Jonathan T. Kolstad
Wharton School, University of Pennsylvania and NBER

January 25, 2013

Abstract

Traditional models of insurance choice, both theoretical and empirical, are predicated on rational choice and risk protection. When these models are taken to data, it is typical to use the choices that consumers make from menus of health insurance options to estimate their risk preferences, conditioning on observed health risk. A key empirical assumption is that risk preferences represent the primary component of persistent unobserved preferences: if other factors such as information about plan options or perceived plan hassle costs also impact choices systematically then risk preference estimates will generally be biased. In addition to having positive implications for choice predictions, omitting such unobserved choice factors can have normative implications for welfare analysis. Such additional factors are generally assumed away because distinguishing between risk preferences and other sources of unobserved heterogeneity is very difficult with administrative data on choices and claims alone.

In this paper we combine administrative data on health plan choices with unique survey data on consumer beliefs and other unobserved preference factors to separately identify risk preferences, various information frictions, and plan hassle costs. These data sets are linked at the individual-level and are used to construct observed measures of information frictions and hassle costs. We develop a simple empirical framework that allows us to account for these additional factors in choice and demonstrate that accounting for these typically unobserved features meaningfully alters estimates for conventional parameters in insurance choice such as risk preferences. We develop a welfare framework that integrates both information frictions and hassle costs and assess the welfare impact of a counterfactual menu design with only a high-deductible health plan option. We illustrate that the welfare implications, and subsequent policy decisions, are quite different when these additional choice factors are accounted for.

*We thank Microsoft Research for their support of this work. We also thank seminar participants at Boston College, Stanford IOFest, USC, the University of Chicago, the University of Rochester, the University of Texas and the ASSA Annual meeting. All errors are our own.

1 Introduction

In both employer-sponsored health insurance markets and recently legislated health insurance exchanges, decision-makers grapple with issues of how to optimally select or regulate plan offerings and prices. How consumers value different health insurance plan attributes is fundamental to health insurance market design. Without knowledge of consumer preferences, it is difficult to appropriately answer key questions such as which type of plans to allow and how those plans should be priced to third-party payers and consumers.

Accordingly, recent empirical work has sought to estimate structural models of consumer choice resulting from specific market designs (see e.g. Cohen and Einav (2007),Carlin and Town (2009),Einav et al. (2010b), Einav et al. (2011), Ericson and Starc (2012), Kowalski (2012), Bundorf et al. (2010) and Handel (2012)). One common aspect in all of these papers is that they rely on detailed administrative plan choice and claims data to identify micro-foundations such as risk preferences (or related phenomena) and consumer expenditure risk. One limitation of using only administrative choice, claims, and demographic data is that the choices consumers make are usually the only instrument available to identify unobserved preference factors, such as risk preferences; the combination of the model applied and the assumptions about the choice process identify the underlying parameters of consumer utility.

There are many potential unobserved preference elements besides risk preferences that can impact demand for distinct insurance plans. Given that health insurance plans are complex financial objects, it is likely that many consumers are not fully informed about different key plan design aspects or even their own medical expenditure risk.[1] Further, time and hassle costs related to actually using an insurance plan differentiate even actuarially identical options.[2] Lack of information about factors relevant to plan utility — choice features we refer to as information frictions — and potential time and hassle costs are unobserved factors that can impact actual plan choices in the same manner as unobserved risk preferences.

Despite the potential importance of these features of plan choice, with the administrative data

[1] Kling et al. (2011) studies an experimental design where certain consumers in Medicare Part D are randomly given additional information about plan features and how they should value them given their past expenditures. They show that information provision has treatment effect that moves a non-negligible group of consumers into plans which provide better value for them, indicating that they were not previously fully informed. Similarly, Abaluck and Gruber (2011) demonstrates systematic choice errors by Medicare Part D beneficiaries in making plan choices.

[2] Making plan operations more efficient and enhancing consumers' ease of use are central motivations for the health insurance exchanges set up in the Affordable Care Act.

sets using in prior studies, it is very difficult to credibly identify the impacts of information frictions and hassle from risk preferences using only consumer plan choices. As a result, to the extent that these additional unobserved preference components are present in the environments studied, they are likely to bias risk preference estimates and the subsequent conclusions drawn about key market design questions. Distinguishing between these choice factors matters crucially for welfare analysis; while risk preferences impact consumer choices as well as consumer welfare once a consumer is enrolled in a given plan, information frictions may impact consumer choices but not consumer welfare once enrolled.[3]

In this paper we leverage new proprietary data from a large firm with over 50,000 employees to separately identify consumer risk preferences from various information frictions as well as from time/hassle costs. Our approach combines the type of detailed administrative data common to the literature with a comprehensive economically motivated survey where consumers' answers are linked to the administrative data at the individual level. The administrative data we collect is a detailed individual-level panel of consumer insurance plan choices, subsequent medical claims, demographics, and employment characteristics. The survey, administered to a random sample of 4,500 employees, asks consumers questions designed to measure various choice frictions including time/hassle costs of plan administration and information frictions relating to financial plan characteristics ((e.g. deductible, co-insurance, OOP maximum), non-financial plan characteristics (e.g. provider network differences), own total medical expenditures, and several other choice-relevant factors.

We illustrate the importance of accounting explicitly for these additional frictions by estimating a series of structural choice models including (i) a baseline model with just risk preferences and health risk (ii) the baseline model incorporating consumer inertia and (iii) our primary model adding information friction and hassle cost measures to the inertial baseline model. The primary model reflects an expected utility maximizing risk averse consumer who may not have full information on available plan attributes/options and may perceive that certain plans have higher hassle costs than others. Comparison between the first two of these models, which are similar to those in the literature, and our full model provides direct measures of (i) the importance of distinct frictions

[3]For example, if a consumer chooses Plan A over Plan B only because they have much more information on Plan A, it is not necessarily true that they would be worse off if Plan A were removed from the choice set and the consumer was forced to enroll in Plan B. In concurrent work, Baicker et al. (2012) studies price elasticity of demand for health insurance with a welfare model that also implies a gap between the choices consumers make and the choices that maximize their welfare if fully informed.

in consumer choices and (ii) how much risk preference estimates are biased by omitting these additional factors. We illustrate the welfare implications of these results by studying the impact of a counterfactual plan menu design that removes all plan options except for a high-deductible health plan (HDHP) that consumers have limited information about.

Prior to estimating our choice model, simple analysis of the survey responses demonstrates that there are important information frictions in the market. Only 27% of the population can correctly identify the defining feature of the HDHP option: the deductible. When asked whether there were any differences in the provider network between the PPO option and the HDHP option only 35% answered correctly that they were identical. When we decompose the responses across the actual plan choices made, we find evidence that information frictions are not only rife in the choice environment but that they seem to be correlated with actual choices made. For example, when asked to identify the deductible in the HDHP option a majority of those enrolled in the plan were able to do so correctly. Among those who remained in the PPO option, only 22% were able to do so. Of those who newly enrolled in the HDHP plan 49% were able to (correctly) answer that the provider networks were identical for the two options. Among PPO enrollees, only 32% answered correctly (14% answered that the PPO had a larger network).

Estimating our choice model with and without inertia and information frictions demonstrates the important role that each plays in preference estimation. Using the baseline of insurance choice that assume rational choice of risk averse consumers, we find a mean coefficient of risk aversion in the population of $5.03 \cdot 10^{-4}$.[4] This level of risk aversion corresponds to a consumer who would only be indifferent between not taking a gamble and a gamble in which he gains \$1000 with a 50% chance and loses \$662 with a 50% chance. He would have to be paid roughly \$340 in expectation to take on a risky bet. While not unprecedented, these results suggest a very high level of risk aversion if we take choices as revealing underlying preferences. Incorporating information frictions into the model, we find that consumers remain risk averse but the magnitude is diminished substantially. Accounting for lack of information, the same consumers would actually be willing to take a gamble with a 50% chance of gaining \$1000 if they lost only \$853 with 50% probability. Taking on the risky bet requires a payment of \$147 in expectation, substantially less than the estimate in the

[4]This model estimates risk preferences conditional on a cost model that predicts a distribution of out-of-pocket expenditures for each family, plan, and time-period. This cost model leverages previous expenditures and the medical implications of previous claims to predict future expenditures in a sophisticated manner.

standard model.

Modeling the welfare impact of moving all individuals into the HDHP (precisely the actual policy taken by the employer we study in 2013) demonstrates the importance of distinguishing information frictions from welfare relevant elements of choice. We estimate the population mean compensating variation due to the switch in the base model (without any inertia or information frictions) is -$1236. Accounting for information frictions and inertia, our model predicts a population mean compensating variation of -$588, roughly half of the welfare loss predicted under the standard model.

Using our model we also study a fundamental trade-off in designing optimal insurance coverage: balancing the welfare gains from risk protection (increasing in generosity of coverage) against welfare losses due to moral hazard (also increasing in generosity of coverage). If information frictions impact plan choices relying on revealed choices in computing the optimal coverage level will lead to an incorrect weight, either on moral hazard or risk protection. For the population who remained in the PPO options, we compute the welfare loss that would result from moving them to the HDHP option. We then compute the implied savings at different levels of demand elasticity, allowing us to solve for the breakeven elasticity of demand that makes it optimal to move from the PPO option to the HDHP option. Using only the baseline model, we find that the HDHP option is only optimal for elasticities of demand for health care services that are greater than .42 to .72. In the full model — taking information frictions and hassle/time costs into account — the elasticity necessary to make the shift from the PPO to the HDHP welfare enhancing is substantially less elasticity. Demand that is more elastic than .21 to .36 makes the move to the HDHP welfare enhancing.

The paper proceeds as follows. Section 2 develops a conceptual theory of insurance choice. Section 3 describes the data, empirical setting and survey. Section 4 develops our empirical model of insurance choice. Section 5 presents results. Section 6 analyses counterfactual policy scenarios and section 7 concludes.

2 Foundations of Choice in the Health Insurance Market

2.1 Standard Model

The canonical model of preferences for health insurance is based on a risk averse consumer who would prefer to pay a fixed premium to avoid losses in the bad state of the world in which he becomes sick (see e.g. Arrow (1963) or Rothschild and Stiglitz (1976)). In this simple case, the decision of insurance plan depends merely on the expected out-of-pocket payment under different scenarios and the risk aversion of the purchaser; health insurance is a tool for financial risk protection. We model this as an individual, indexed by k, choosing health insurance plan j. The individual's utility from the choice can be written as:

$$u_{kj} = \int_0^\infty u(W_k - P_{kj} - OOP, \gamma_k) f_{kj}(OOP|\psi_j, \mu_k) dOOP \tag{1}$$

where W_k is wealth, P_{kj} is the premium facing individual k in plan j. $f_{kj}(OOP|\psi_j, \mu_k)$ is the probability density of out-of-pocket expenditures in plan j for individual k. Out-of-pocket spending is determined in each plan by two features: the plan design, indexed by ψ_j and the consumer type, indexed by μ_k, that captures the expected total spending. Together, the terms of the plan and the total spending define the joint density of out-of-pocket spending. The term γ_k is a coefficient of risk aversion for individual k.

This simple framework captures the standard model of preferences for insurance; Individuals are willing to pay a higher premium for a plan if they have a lower expected out-of-pocket spending given their expected expenditures and the more risk averse they are. The individual making a choice in this model has uncertainty over health care expenditures in different states of the world. However, he does know with certainty the density of expenditures — implicitly he is able to place a probability weight on each of the different illnesses that might befall him, know how much the appropriate treatment would cost and understand the terms of the different plan options that result in different rates of cost sharing depending on expenditures/illness states. This workhorse model has a number of important advantages. It is a tractable representation of preferences with a clear empirical analog. The elements of the model can also be observed in widely available data sets (e.g. the expected expenditures for an individual and the plan options).

2.2 Non-Financial Attributes in Plan Choice

To better reflect actual choices, we must account for the fact that modern health insurance is not a purely financial product. With the rise of managed care and alternate benefit designs, the insurance one holds can determine the type of care available, the price paid and the hospitals and doctors one can access. The introduction of Health Savings Accounts (HSA) and Flexible Spending Accounts (FSA) have introduced non-price attributes of plans such as hassle of dealing with reimbursement and billing. More generally, health insurers are differentiated products across a variety of dimensions beyond financial risk protection.

We extend the model to account for additional components of the choice problem that are not directly related to financial risk.[5] Plans differ by the network of physicians and hospitals available, the hassle-cost associated with dealing with claims, the subsidies available for each plan and the tax preference for health insurance benefits. Each plan has a plan specific shifter $\pi_j(\psi_j, \mu_k, (1-t_k))$. We allow the plan shifter to depend on plan design (ψ_j) and consumer type (μ_k) to reflect the fact that hassle costs depend on consumption of care and illness. Similarly, the plan shifter depends on an individuals tax rate as the value of HSA contributions depends on the marginal tax rate Incorporating these features into the model utility from plan j for individual k is:

$$u_{kj} = \int_0^\infty u(W_k - P_{kj} + \pi_j(\psi_j, \mu_k, (1-t_k)) - OOP, \gamma_k) f_{kj}(OOP|\psi_j, \mu_k) dOOP \tag{2}$$

2.3 Information Frictions in Plan Choice

In the description above, the choice of insurance plan relies entirely on individuals risk preferences, their expected expenditures and plan attributes. Importantly, when individuals make insurance choices we assume they can access the necessary to make the correct decision (even if it is a decision made by taking an expectation over uncertain outcomes). Accordingly, individual choices reflect real preferences for trading off premiums in exchange for shifts in expected OOP spending or non-financial attributes across different plans. This assumption is critical and underlies positive analysis of choice patterns as well as assessing welfare from using revealed preference becomes more challenging (Spinnewijn (2012) and Bernheim and Rangel (2009)).

[5]The inclusion of these features in models of insurance choice is not new (see e.g. Ho (2009), Cutler et al. (2000)). Measurement of the plan attributes, however, has typically posed a problem.

There are, of course, many ways that choices could differ from the model described in equation 2. The feature that is perhaps most critical and potentially unlikely to hold in practice is information availability. The optimization problem assumes individuals are able to form an estimate for the likelihood of different illnesses. We assume that the individual is able to forecast the cost of becoming sick with a particular condition. This includes not only the actual treatment but also the uncertainty over price differences across treatment locations and regimes. The other critical feature of the choice environment is that the individual is aware of the benefit design and attributes of the different plan options. Without this, not only would the individual be unaware of attributes that differentiate each plan, but even in the standard model where preferences are merely over financial risk, would not be able to forecast their OOP spending in the different plan options. Similarly, individuals may not have perfect information on the non-financial attributes of plan options (e.g. provider networks and hassle costs), particularly in the absence of having experience with a plan. To model information frictions we allow the true value of the key parameters of the choice model to be observed with error. We can express each feature as:

$$\widehat{\mu_k} = \mu_k + \delta_k^{\mu} + \epsilon_k$$
$$\widehat{\psi_j} = \psi_j + \delta_j^{\psi} + \epsilon_j$$
$$\widehat{t_k} = t_k + \delta_k^{t} + \epsilon_k$$
$$\widehat{\pi_j} = \pi_j + \delta_k^{\pi} + \epsilon_j$$

We assume that individuals observe each plan attributes with two types of error. The first is standard measurement error (mean zero) captured by ϵ. The second is an attribute specific shifter, δ, that captures information frictions in the model. Consumer choice no longer reflect the attributes of the plans (and preferences over those attributes) but, instead, beliefs about those attributes that could be incorrect. Incorporating these features into the choice model, consumers choose an insurance plan based on their beliefs about plan attributes and cost according to:

$$u_{kj} = \int_0^{\infty} u(W_k - P_{kj} + \widehat{\pi}_j(\widehat{\psi_j}, \widehat{\mu_k}, (1 - \widehat{t_k})) - OOP, \widehat{\gamma_k}) f_{kj}(OOP|\widehat{\psi_j}, \widehat{\mu_k}) dOOP \tag{3}$$

From (3) we see how the information frictions can impact the choice behavior of consumers in

important ways. Because both $\widehat{\psi_j}$ and $\widehat{\mu_k}$ enter the choice problem and rescale the response to OOP expenditure risk choice, even if we observe the choices of individuals who truly optimization given their beliefs we cannot recover key features of the model such risk preferences. Similarly, if individuals are imperfectly informed about the non-financial attributes of the plan, $\widehat{\pi_j}(\widehat{\psi_j}, \widehat{\mu_k}, (1 - \widehat{t_k}))$ this will lead to choices that differ from what would have occurred with full information on the plan's network of physicians, true hassle costs or an understanding got the tax preference for benefit features. While choices are affected by them, information frictions do not enter true utility (captured in equation (2)). Thus, when information frictions impact choices, the standard model will predict behavior that differs from practice. Furthermore, taking observed behavior as revealing preferences will lead to biased estimates and, therefore, incorrect assessment of the welfare impact of different counterfactual market or policy scenarios.

Whether information frictions exist in practice and, if so, how important they are is an open question. Addressing this empirically has been a challenge because the data requirements are substantial. To compare the model in equation (2) to equation (3) requires both data on actual choices and plan attributes as well as a measure of beliefs about plan attributes. Our empirical setting provides exactly that, by combining administrative data on claims and choices of insurance with a detailed survey on consumer beliefs about attributes of the plan options. The remainder of the paper focuses on develop a structural empirical model that corresponds to equation (3). This allows us to assess the positive impact of information frictions on choice as well as the impact of information frictions on welfare under different counterfactual scenarios.

3 Data and Empirical Setting

Critical to our ability to model consumer choices incorporating information frictions is the fact that we have detailed administrative data on actual benefits choices made by employees and detailed claims data on health care utilization over time combined with survey evidence that allow us to capture many of the non-standard features that might impact choice. Adding to our ability to model choice is the relatively clear choice environment at the firm we study as well as some of the unique features of the choice. In this section we outline the setting and how choices were made. We also describe the different data sets in detail.

3.1 Background on the Choice Environment

We study benefits choices made by the universe of employees at a single large firm in 2012. Employees choose between two different health insurance options: a PPO option with generous first dollar coverage and HDHP option. Depending on the location of the office, a small share of employees could also choose a Health Maintenance Organization (HMO) option. [6]

The PPO option has the largest share of employees and has been the primary health insurance plan for many years prior to the introduction of the second plan. There is no deductible for the plan, nor is there any cost sharing for health care services consumed within the network of providers. The plan offers a very broad network of primary care doctors, hospitals and specialists. Beginning in 2009, the firm introduced a HDHP option. Our study takes place in a particularly interesting time for the firm. They are transitioning away from the traditional PPO options towards a HDHP option. As of 2013, all employees are enrolled in the HDHP.

Table 3.1 compares the two plans of interest — the PPO and the HDHP. They differed substantially in the benefit design features that determined OOP cost and the financial incentives to adopt them (e.g. subsidies and tax preferences for savings in the HSA account associated with the HDHP). However, because they were administered by the same insurer and were self-insured plans, they were identical in all other key features. For example, the HDHP had an identical network of providers as the PPO plan. The HDHP plan has a substantial deductible, $1500 for individuals, $3000 for a couple and $3750 for a family. Once the employee spends an amount in excess of the deductible he must then pay co-insurance of 10% for providers in-network and 30% for providers outside of the network until his total spending exceeds the out-of-pocket maximum — $2500 —- at which point all expenditures are paid by the insurer. The plan is also linked to a HSA that allows employees to contribute pre-tax up to a maximum of $3200 for individuals and $6250 for families. The employee can then use money from their HSA to pay for any qualifying health expenses without paying a penalty. They can also, however, save this money to be withdrawn at a later date either for use on health expenditures, in which case they do not pay any penalty, or to be used on other consumption in which case they must pay tax on the savings in the account at their *current* tax rate and a small additional penalty. Employees do not pay a premium for either plan so there

[6] Relatively few actually chose this option – approximately 5% – and this group remained steady over our sample period (2009-2012).

Health Plan Characteristics		
	PPO	HDHP
Premium	0	0
Health Savings Account (HSA)	No	Yes
HSA Subsidy*	-	$1,500
Max. HSA Contribution**	-	$3,100
Deductible*	0***	$1,500
Coinsurance (IN)	0%	10%
Coinsurance (OUT)	20%	30%
Out-of-Pocket Max.*	0***	$2,500
Provider Network	Same as HDHP	Same as PPO

* Employee+1 Tier is 2x Individual Amount, 3+ Tier is 2.5x Individual Amount
**For 2 or more, $6,250 is max. contribution
***For out-of-network spending, PPO has a deductible of $100 per person (up to $300) and an out-of-pocket max. of $400 per person (up to $1200)

Table 1: This table presents key characteristics of the two primary plans offered at the firm we study. The PPO option has more comprehensive risk coverage while the $HDHP$ option gives a lump sum payment to employees up front but has a lower degree of risk protection.

is no difference in initial price between plans.

In Figure 3.1 we plot the share of enrollees in each of the different plans at the employer from 2009-2012. The share of employees choosing the HDHP plan remained below 1% for the first two years it was offered. Beginning in 2011, motivated in part by the coming complete transition to the HDHP plan, the employer began an extensive education and marketing campaign to inform employees about plan options and to facilitate movement to the HDHP. In 2011 there was a substantial uptick in enrollment to 8.26%, potentially resulting from the information campaign. By 2012, the share in the HDHP option had increased to 13.25%. The increase in HDHP enrollees appears to be drawn from the traditional PPO plan as the share in other plan options remained similar in 2011 and 2012 as the HDHP enrollment increased.

Despite the enrollment growth, we still see a relatively small share of employees choosing the HDHP plan. To provide some preliminary evidence on the choice process, we can compare the rates of HDHP adoption with the *ex post* optimal plan choice based on each members actual

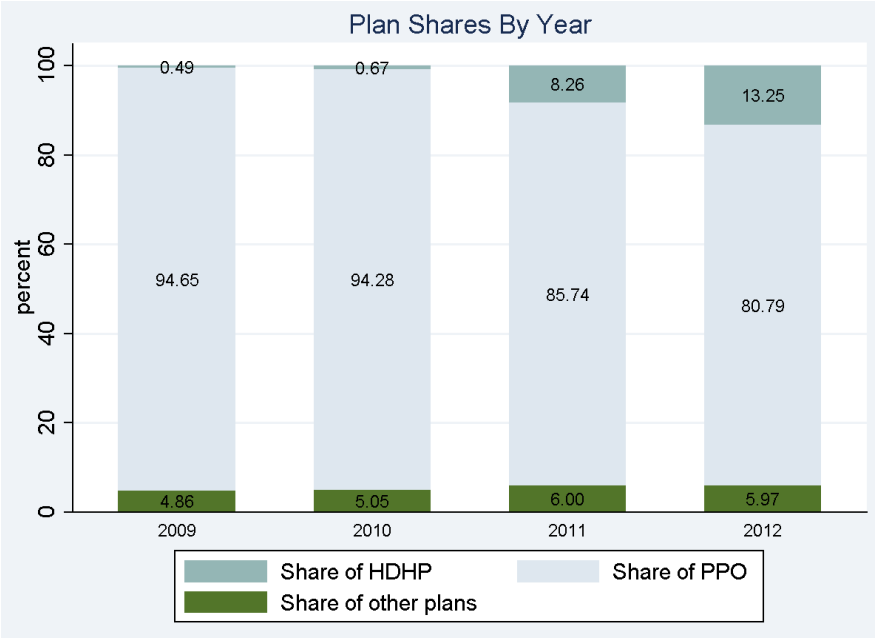


Figure 1: Share of Employees in Each Health Insurance Option by Year

health expenditures. This simplifies the choice process and ignores the primary benefit of health insurance: risk protection. Nevertheless, it sheds some light on the potential magnitude of choice errors that could be occurring. Figure 3.1 depicts the financial returns to selecting the HDHP option relative to the PPO option for a representative individual employee.[7] The x-axis plots realized health expenditures and the y-axis plots the according financial returns for the HDHP relative to the PPO. At low levels of health care cost, the HDHP plan provides a relatively large financial return. Employees do not pay a premium for either plan but they do receive a substantial contribution to their HSA from the employer. In addition, the ability to contribute to their HSA account pre-tax also provides financial returns. As spending increases, these gains are reduced dollar-for-dollar below the deductible level. Once out-of-pocket spending surpasses the deductible they pay coinsurance that diminishes the slope of the loss in the HDHP relative to the PPO as spending increases. Finally, beyond the stop loss level the employee pays no additional amount out-of-pocket. The graphical depictions demonstrates that there is a unique level of expenditure at which the PPO plan is *ex post* preferred to the HDHP plan. Furthermore, the maximum financial loss from choosing the HDHP plan is actually relatively low, even for very high levels of expenditures ($500).

[7]The same benefit structure holds for couples and families with shifts in the levels of the key plan terms.

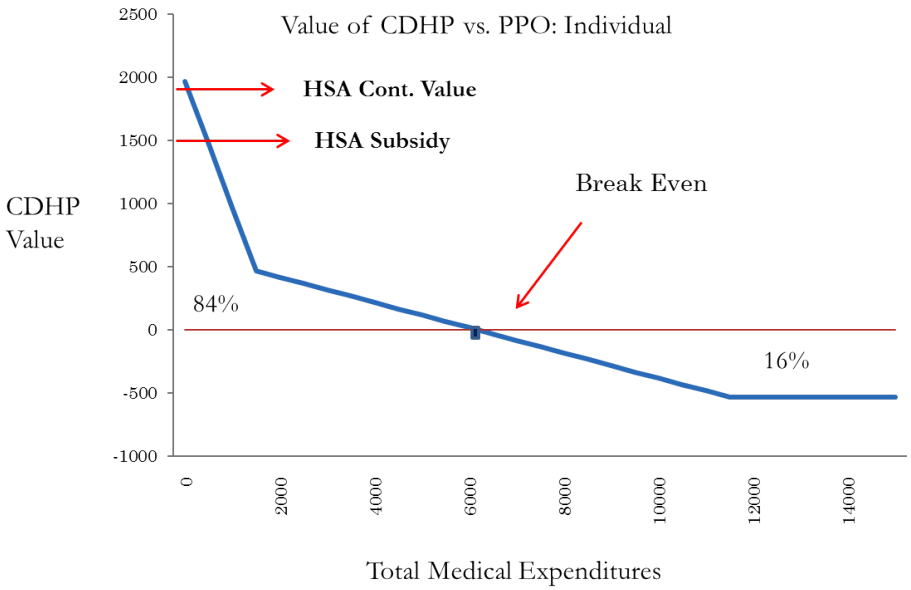


Figure 2: Break Even Analysis for Employees Covering Only Themselves, 2012

Extending the analysis in Figure 3.1 we can compute the share of employees whose total medical expenditures were below the break-even point in 2011 for individuals, couples and families — those who *ex post* were better off in the HDHP. Assuming all employees contributed the maximum amount to their HSA 73% of employees should have switched in 2011 to the HDHP. With only 50% or zero contribution to the HSA, 60% and 35% of employees would have been better off financially in the HDHP plan in 2011. Yet, as is clear in Figure 3.1, only 8.25% actually made that choice in 2011 and only 13.25% by 2012.

This simple comparison suggests there is a more complex choice process occurring than merely comparing financial outcomes. The remained of the paper focuses on understanding what additional features of the model can explain this observed behavior. In the standard model of insurance choice, the observed choices can be rationalized by particularly risk-averse consumers; they are willing to leave large gains on the table if they do not get sick in order to avoid the possible loss of $500 were they to get sick. Employee behavior could, however, result from lack of information on plan features, likely expenditures or beliefs about non-financial attributes of the plan (i.e. hassle-costs, physician networks, etc.). To capture these features we developed a survey instrument to ask a subset of employees about these features directly.

3.2 Administrative Data on Choices and Claims

One of our key data sources is claims data from the large employer we study. The employer is a large employer with 41,361 employees in the United States. This corresponds to approximately 102,000 covered lives in the health plans we study. Our administrative data contain three key components. We see the choice of health plan facing each employee and the health plan they enrolled in. For each employee who enrolled in the HDHP, we also observe contributions to the HSA made by the employer and the employee. In addition, we observe the universe of health care claims for all covered individuals in all plans. These data are available from 2009-2012. The first column of Table 3.2 presents summary statistics for the full sample. One thing that is clear from the summary statistics is that the population we study is not representative of the general population. The employee population is heavily male (76.4%), relatively young and has relatively high income. While this makes the study less representative, we are particularly interested in the results insofar as this population seems more likely to have the education, resources and cognitive skills to overcome information frictions.

3.3 Survey Data and Design

In order to measure information frictions and beliefs about non-financial plan attributes, we developed a survey instrument. Our survey instrument was designed in conjunction with the Human Resources department at the employer we study. The survey was released in March of 2012 and was sent to 1,500 employees in each of three groups (totaling 4,500): employees who had been in the HDHP for the previous year and remained in it ("incumbents"), new HDHP enrollees in 2012 and those in the PPO plan. Of the 1,500 initially contacted in each group, we received response from 579 incumbent HDHP enrollees, 571 new HDHP enrollees and 511 PPO enrollees.

The average response rate overall was 38%. The last two columns of Table 3.2 compare summary statistics for the targeted population for the survey and the actual respondents to the survey (reweighed to reflect the actual plan choice composition in the market) with the full sample. The populations are, on the whole, quite similar. We note that the targeted survey population, selected at random by the employers HR department, appears to be somewhat younger, have a higher share with incomes in excess of $125,000, have fewer families and slightly lower health care spending. However, when we turn to the third column, we actually see that the respondents whose survey

Sample Demographics	Full Sample	Survey Rec. Resampled	Survey Res. Resampled
No. Employees-Actual	41,361	4500	1661
No. Employees-Resampled	-	18000	6260
2011 PPO%	88.8	89.6	88.7
2012 PPO%	82.7	83.0	81.6
Gender (% Male)	76.4	76.8	75.6
Age			
18-29	8.6%	14.9%	11.6%
30-39	41.1%	43.8%	42.7%
40-49	38.1%	32.7%	34.1%
50-59	10.9%	7.7%	10.5%
≥60	1.3%	0.9%	1.2%
Income			
Tier 1 (< $75K)	2.7%	2.2%	2.2%
Tier 2 ($75K-$100K)	10.1%	13.1%	14.0%
Tier 3 ($100K-$125K)	35.3%	38.9%	37.9%
Tier 4 ($125K-$150K)	30.5%	29.6%	31.3%
Tier 5 ($150K-$175K)	12.0%	10.8%	10.0%
Tier 6 ($175K-$200K)	4.7%	3.5%	2.9%
Tier 7 ($200K-$225K)	2.0%	1.0%	0.9%
Tier 8 ($225K-$250K)	0.7%	0.1%	0.0%
Tier 9 (> $250K)	0.8%	0.1%	0.0%
Family Size			
1	23.0%	29.0%	20.9%
2	19.0%	19.4%	21.9%
3+	58.0%	51.6%	57.2%
Family Spending			
Mean	$10,191	$8,820	$11,247
Median	$4,275	$3,363	$4,305
25th	$1,214	$878	$1,176
75th	$10,948	$9,388	$11,555
95th	$35,139	$32,171	$41,864
99th	$87,709	$80,370	$87,022

Table 2: This table presents summary demographic statistics for the samples we study. The first column represents all employees who were present in our data and have complete records for at least eight months in 2009, 2010, and 2011, and the first month of 2012. The second column represents all employees who received our survey, regardless of whether or not they responded. The third column represents all employees who responded to our survey. Statistics from gender onwards represent only 2011, and use the resampled statistics for the second and third columns.

responses we use in our empirical models, are quite similar to the overall population across all of these measures. The average spending is slightly higher among the respondents compared to the overall population but comparing spending at different points in the distribution this appears to be driven by higher spending in the upper tail of cost and not a systematically higher spending group of respondents than the overall population.

Each multiple choice question was based on economic theory underlying information frictions and plan choice. Table A in the appendix summarizes the questions used in our analysis and the responses from the survey population. The questions focus on four major areas of the benefits choice. The first targeted area assesses knowledge of the benefit design in the HDHP plan. These questions target information frictions directly as they ask respondents to answer key features of the HDHP plan. Each respondent was asked to identify the deductible, coinsurance rate, out-of-pocket maximum, HSA subsidy level and tax benefits for HSA contributions out of a set of choices. The second set of questions focused on a related sources of information frictions, beliefs about plan attributes and medical expenditures. Respondents were asked whether the PPO or HDHP had any differences in the networks of providers available through each (recall they are identical). We also ask a set of questions of respondents to determine whether they were able to assess past medical expenditures and likely future medical expenditures. The third area of focus was on time and hassle costs associated with the HDHP plan. These include the time and resources required to manage their HSA and HDHP (e.g. collecting and submitting receipts for care to be reimbursed from their HSA). In addition to directly eliciting beliefs about the time required, we also ask questions about preferences for hassle in the HDHP. Finally, we ask a set of questions to ascertain the amount of effort that went into an employees choice, the clarity of their beliefs about the plans and their satisfaction with their choice.

3.4 Survey Evidence

Before turning to our formal model of choice that incorporates survey data, we present some simple results from the survey. These raw results demonstrate the power of incorporating survey data with administrative data on actual choices and the potential importance information frictions. There are clear patterns in the raw survey responses that are consistent with limited information among those choosing plans. Furthermore, answers to survey questions have a strong gradient with respect

to actual plan choices made.

For example, Figure 3.4 plots the share of individuals who correctly identified the deductible in the HDHP plan by the by their previous plan choices. A (slim) majority of employees who were enrolled in the HDHP were able to correctly answer the question. On the other hand, only slight more than 20% of employees who remained in the PPO could answer the question correctly. In fact, more PPO enrollees answered incorrectly than correctly, though the majority were "note sure." This plot is indicative of the answers we see across our survey results. In particular, across all of our survey responses we see a strong gradient in the answers with HDHP enrollees displaying substantially more knowledge of the benefit, though we also find a substantial portion of the population is unable to answer the question (answers "Not Sure") in every one of our survey questions. Figures 3.4 and 3.4 present the results for other questions on benefit design in the HDHP.

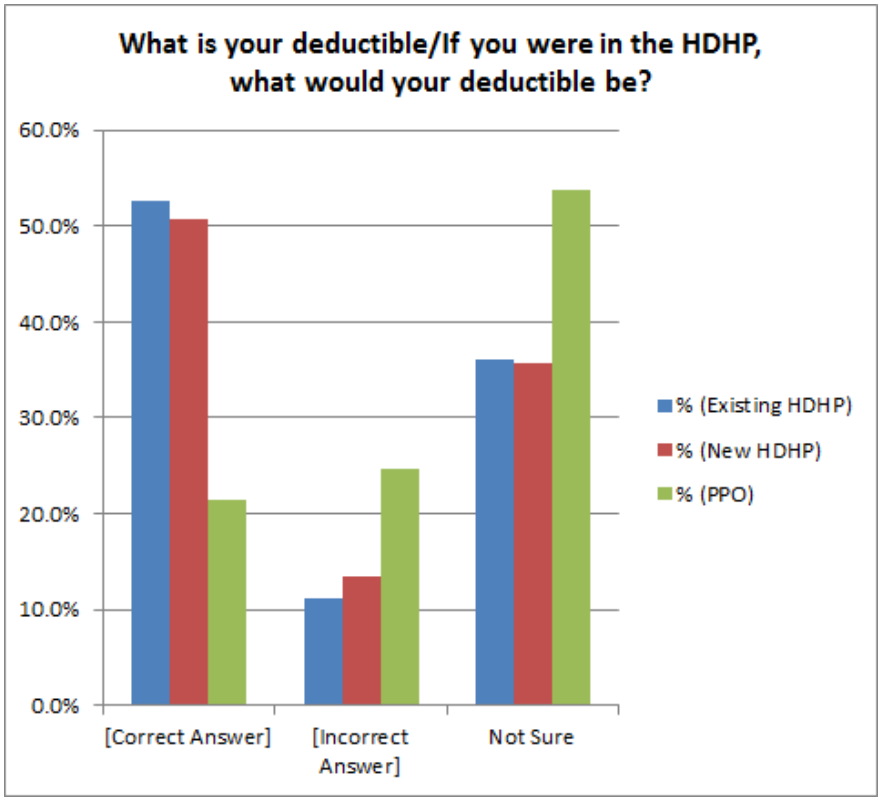


Figure 3: Knowledge of Deductible by 2012 Plan Choice

In order to accurately, or even reasonably, forecast expenditures in the coming year individuals likely need information on how much the spent in the previous year. Figure 3.4 plots survey

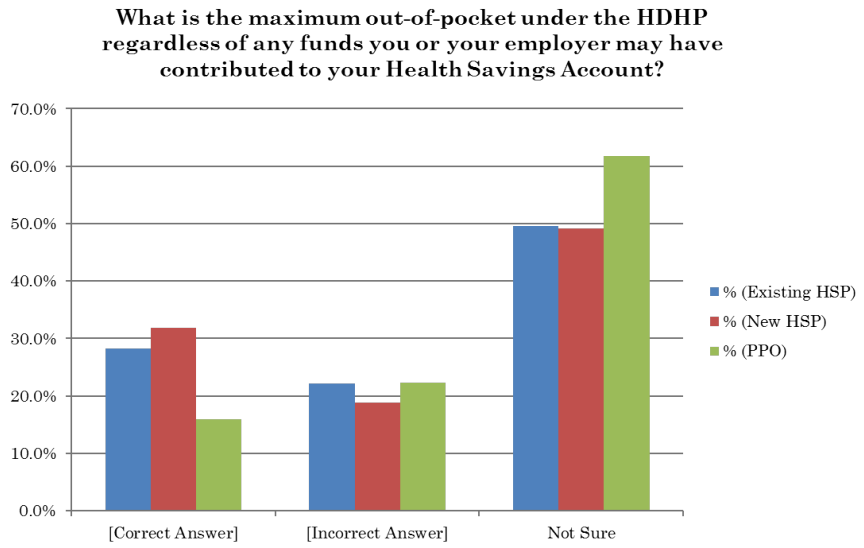


Figure 4: Knowledge of Out-of-Pocket Maximum by 2012 Plan Choice

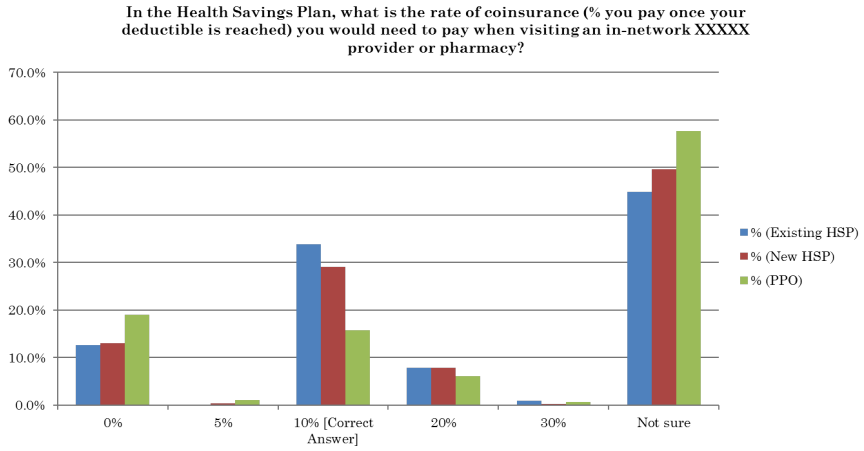


Figure 5: Knowledge of Coinsurance by 2012 Plan Choice

responses to a question asking individuals to forecast the total cost of medical care consumed by them and their dependents between January and December of 2011. The question asked the respondent not to precisely name the total, but to select between $0-500, $501-2,500, $2,501-5,000, $5,000-10,000 and more than $10,000. We see a steep gradient in the likelihood of being able to correctly assess expenditures, with incumbent HDHP enrollees most likely to be correct. Only between 35% and 25% of employees answer the question correctly. Interestingly, it appears a larger share of all of the plan groups underestimates their total medical expenditures than overestimates, though combining these two answers accounts for around 60% of all of the groups. When we subsequently asked survey respondents to provide their confidence in their estimate of their past

year total medical expenditures we find that the majority of the population of both incumbent and new HDHP enrollees and PPO enrollees reply they are somewhat or very confident in their estimate. There is also not a clear gradient in this reported confidence across plan types. Thus, it not only appears individuals are not well equipped to estimate their total expenditures in the past year, even to the level of expenditure buckets, people do not appear to internalize this uncertainty.

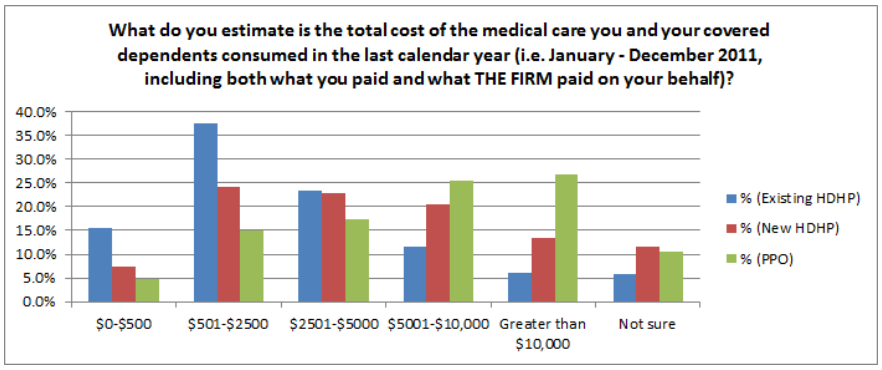


Figure 6: Accuracy of Knowledge of 2011 Total Medical Expenditures

We also ask all survey respondents whether the network of providers is larger, smaller or the same in the HDHP plan. Recall that the network is, in fact, identical. We find that [z] . This level of incorrect and uncertain beliefs about a plan attribute that was both relatively straightforward to consider and emphasized in the information provided by the employer underscore the role of information frictions. To get an idea of how important these information frictions are in explaining choices, we compare the actual 2012 choices to the *ex post* optimal choice based on 2011 data. The results are presented in Figure 3.4. The left panel presents the share of enrollees in the HDHP and PPO plan respectively by their answer to the survey question on provider network. It is clear that those who understood the network was the same were far more likely to select the HDHP plan (23% compared to 6% among those reporting the PPO had a larger network and 12% among those reporting the HDHP had a larger network). The right panel presents the optimal choice based on actual expenditures (ignoring risk preferences). The predicted choices are higher across all survey responses. However, the difference among those who answered correctly are far smaller (23% in the HDHP in practice versus 33% predicted). Among those answering incorrectly, however, the differences are very large. Respondents who believed the PPO had a larger provider network only chose the HDHP plan 6% of the time even though they would have been better off in the HDHP plan 40% of the time. A similar gap exists for individuals who though the HDHP had a

better network (12% actual versus 24% predicted) and who were "not sure" (17% actual versus 39% predicted).

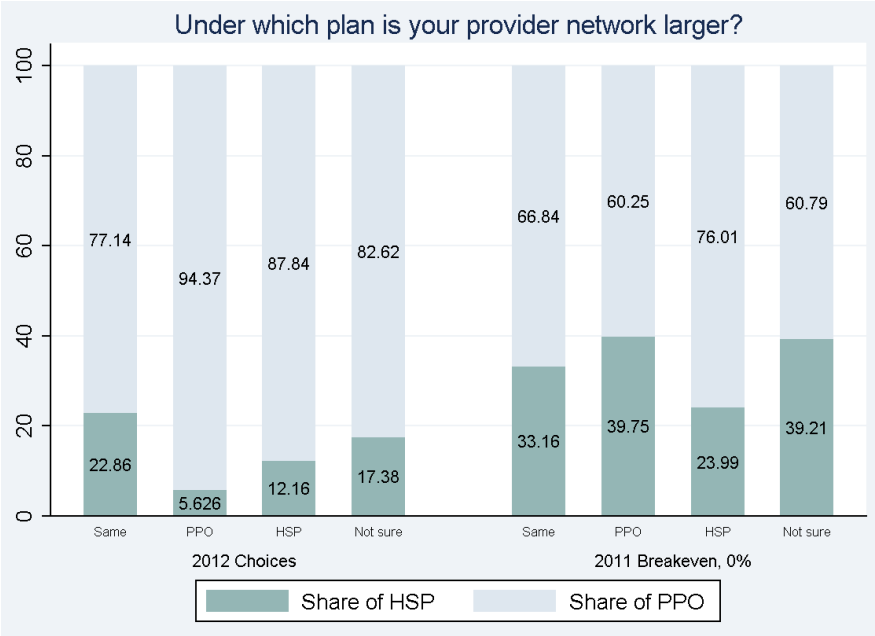


Figure 7: Actual versus Predicted Plan Choices by Knowledge of Plan Provider Networks

In addition to information frictions, our survey data allows us to assess plan attributes that are generally hard to measure. The hassle of dealing with paying for medical expenditures directly and being reimbursed is a potentially important non-financial attribute of the HDHP plan the might impact choice. [8] In order to better understand employees beliefs about those costs, we asked survey respondents about their perception of the time they would need to spend dealing with the HDHP plan as well as their preferences for spending that time. That is, do they "understand" the need to spend some time managing their expenses in the HDHP. These responses are presented in Figures 3.4 and 3.4 below. The results point to a substantial difference in perception of the time required to deal with the HDHP among those who had experienced it and the PPO respondents. A far larger share of PPO respondents reply they believe HDHP enrollment will require in excess of 20 hours to deal with. It is also interesting that the new HDHP enrollees appear to have quite similar beliefs about the time required as those who have actually experienced the plan, the incumbent HDHP enrollees. Figure 3.4 also show a strong gradient with respect to how accepting HDHP enrollees are of the time required to deal with the plan. Only 10% of PPO enrollees report not being concerned

[8] This concern was discussed widely in the focus groups we conducted with employees.

that they may need to spend time managing health care costs.

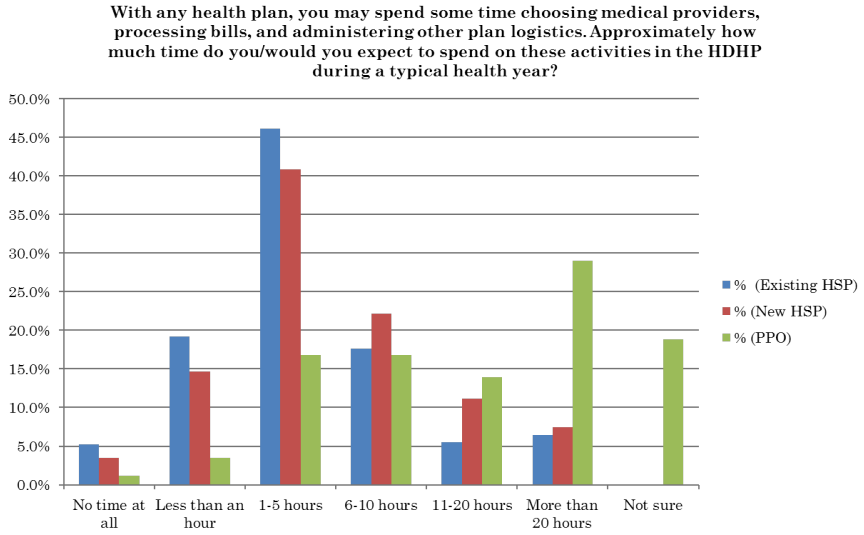


Figure 8: Beliefs about Amount of Time Necessary to Deal with the HDHP Plan

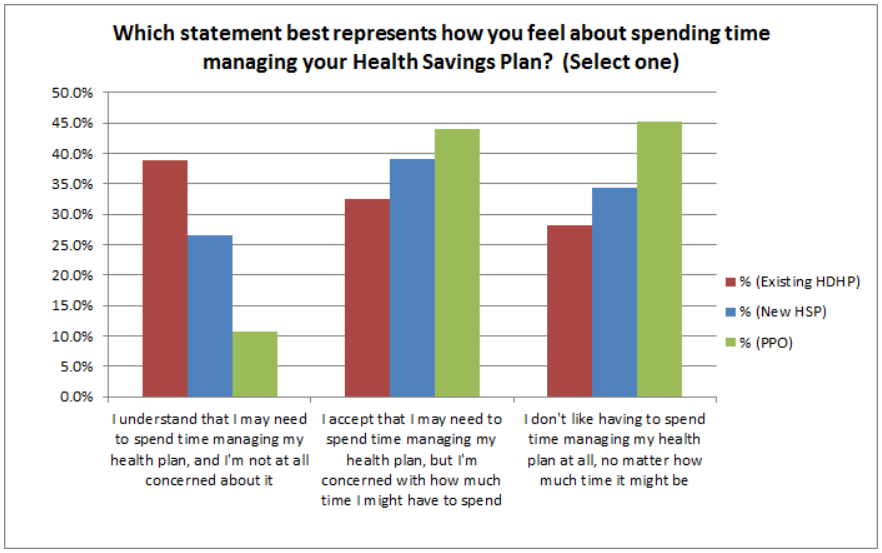


Figure 9: Preferences on Hassle Cost Associated with HDHP

4 Empirical Framework

The analysis in the previous section provides robust evidence that information frictions (i) are present for a variety of key choice dimensions and (ii) are correlated with the consumers' health plan choices. Prior model-driven work on consumer health plan choice has generally been unable to

account for either information frictions or perceptions of time and hassle costs. As a result, estimates of risk preferences in past work may in part account for the impacts of these factors on choices and either over or underestimate true risk preferences. This matters for economic analyses for several reasons. First, we may inherently care about risk preference estimates to better understand consumer choices under uncertainty. Second, accounting for micro-foundations in a more nuanced way allows for more precise prediction of counterfactual choices. Third, the welfare implications of risk preferences for specific policies or economic environments are generically different than those from either information frictions or time and hassle costs.

We develop a series of models to explore these issues and illustrate the impact of accounting for these additional micro-foundations in health plan choices. Our empirical implementations of these models directly link our plan choice, medical cost, and information friction survey data to these underlying economic choice fundamentals. The conclusions we draw from these models should be viewed in the context of the structural assumptions included.

Baseline Choice Model. We first develop a baseline model with two primary components (i) risk preferences and (ii) *ex ante* cost projections. This model is meant to be a basic version of the choice models used in the insurance choice literature. In this section we describe the baseline choice framework *conditional* on our ex ante cost projections: following our discussion of all the choice models studied we describe the detailed cost model that generates these expenditure distributions (this model does not vary with the choice specification).

Denote the family-plan-time specific distributions of out-of-pocket health expenditures output by the cost model as $F_{kjt}(\cdot)$. Here, k is a family unit, j is either the PPO or $HDHP$ plan available at the firm, and t is one of two years from t_0 to t_1. 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:

$$U_{kjt} = \int_0^\infty f_{kjt}(s) u_k(W_k, x, P_{kjt}) ds$$

Here, $u_k(\cdot)$ is the v-NM utility index and s is a realization of out-of-pocket medical expenses from $F_{kjt}(\cdot)$. W_k denotes family-specific wealth. P_{kjt} is the family-time specific premium for plan j. In our setting, as described in more detail in section 3, the premium is 0 for the PPO option and negative, in a traditional sense for the $HDHP$ option, since employees are given money to enroll in that plan. Formally, in our setting we define the premium P_{kjt} for the $HDHP$ plan as:

$$P_{k,HDHP,t} = -(HSA^S_{kt} + \tau_{kt} HSA^C_{kt})$$

Here, HSA^S_{kt} is the firm's subsidy to each employee's health savings account (HSA) when they sign up for the HDHP plan. This is deterministic conditional on the number of dependents being covered as discussed in section 3. $\tau_{kt} HSA^C_{kt}$ is the value of incremental contributions consumers make to the HSA (which they can only make when signed up for the $HDHP$ plan.) This equals the incremental contribution made above the subsidy, multiplied by a family's marginal tax rate.[9] Empirically, we model HSA^C_{kt} based on actual contributions made by those who have signed up for the $HDHP$ plan in the data. This model yields a family-specific prediction of incremental HSA^C_{kt}, denoted $\widehat{HSA^C_{kt}}$, which is inserted into the model such that $P_{k,HDHP,t} = HSA^S_{kt} + \tau_{kt}\widehat{HSA^C_{kt}}$. Appendix B discusses the model for HSA^C_{kt} in more detail.

Given this setup, we assume that families have constant absolute risk aversion (CARA) preferences implying that for a given ex post consumption level x:

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

Here, γ_k is a family-specific risk preference parameter that is known to the family but unobserved to the econometrician. We model this as a function of employee demographics X^A_k. 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.[10]

In our baseline empirical specification a family's overall level of consumption x conditional on

[9] Incremental contributions to the HSA have value equal to $\tau_{kt} HSA^C_{kt}$ if at any point in the employee's life their family spends that money on health care. If they spend part or none of those incremental funds on health, then the value of these incremental contributions is lower. We do not incorporate the value of the $HDHP$ as a tax-free investment vehicle, which would increase it's value depending on the assumed ROI.

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

a draw s from $F_{kjt}(\cdot)$ is:

$$x_{kjt} = W_k - P_{kjt} - s + \epsilon_{kjt}$$

Here, ϵ_{kjt} is an family-plan-time specific idiosyncratic preference shock. Since the plans we study are only differentiated by financial characteristics and the HSA (and not providers) we also follow Einav et al. (2011) and study a robustness check with no idiosyncratic preference shock. Subject to this model, families maximize choose the plan j each year that maximizes U_{kjt}.

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 later in this section. This assumption could be incorrect for at least two reasons. First, families may have private information about their health statuses that is not captured in the detailed prior claims data. We study a robustness analysis where we show that small deviations in our estimates for F_{kjt} do not markedly impact other parameter estimates so only consumers with substantial private information can bias our estimates. Given our detailed individual-level claims data, we believe it is unlikely that there are many consumers with substantial private information in our data. Second, families may have *less* information about their projected future health expenditures: the cost model utilizes a full profile of past claims data in conjunction with medical software that maps past claims to future expected expenses. Further, the baseline model contains the assumption that consumers have full knowledge of health plan characteristics and incorporate that knowledge into their decision process. Our full model, which incorporates our detailed individually-linked survey data about plan and health risk knowledge, addresses a variety of ways in which consumers have limited information about potential out-of-pocket expenditures when choosing a health plan.[11]

Inertial Baseline Model. The baseline model presented above reflects consumers making active plan choices in an environment without information frictions or additional micro-foundations like time or hassle costs. In this section we add inertia to the baseline model while in the next we add information frictions and time and hassle costs.

In the choice environment that we study, existing employees have a default option in their

[11]To the extent that we cannot fully account for limited information about F_{kjt} with our survey questions, our robustness analysis that examines the impact of small deviations in F_{kjt} illustrates robustness with respect to small deviations in perceptions about expense risk from our baseline model.

yearly health plan choice: if they take no action at all they are enrolled in the plan they previously chose. As shown in prior work (see e.g. Handel (2012)), in a default option environment, inertia, defined as choice persistence not resulting from stable preferences, can have a substantial impact on the choices that are made and, consequently, on consumer welfare. Here, we add inertia to the baseline model as an implied monetary cost of switching plans when a default option is present, similar in structural interpretation to a tangible switching cost. Inertia changes the baseline model by augmenting consumption x_{kjt} as follows:

$$x_{kjt} = W_{kt} - P_{kjt} - s + \eta(X_k^B)\mathbf{1}_{j_t=j_{t-1}} + \epsilon_{kjt}$$

Here, η represents inertia and depends on observed demographic variables X_{kt}^B, which are described in more detail in the estimation section. Apart from the inclusion of η the model with inertia is identical to the baseline model: consumers have CARA risk preferences and choose the plan that maximizes their utility U_{kjt} subject to their out-of-pocket risk F_{kjt}.

There are several implicit assumptions in the inertial model, in addition to those in the baseline model, that warrant discussion. First, inertia is modeled in a specific way, as an incremental cost paid conditional on switching plans. This framework implies that, on average, for a family to switch at t they must prefer an alternative option by $\$\eta$ more than their default. This follows the approach used in Handel (2012) as well as other prior work modeling inertia and switching costs (see, e.g.Shum (2004) or Dube et al. (2008)). There are multiple potential underlying micro-foundations for inertia, each of which could correspond to an alternative model. Handel (2012) describes these potential models in detail and discusses why differentiating between them may not be important to answer certain economic questions. In our setting, we identify the extent of inertia by comparing the relative value of health plan choices made by new employees, who make active plan choices with no default option, to similar looking existing employees who do have a default option. Given this identification, it is unlikely that our specific representation of inertia impacts estimates of other preference parameters (such as risk preferences) because those parameters are identified separately from inertia, in any form, by new employee active choices.[12] We present a

[12]Our counterfactual menu design analysis assumes a forced or active choice environment so as long as non-inertial preferences are unbiased (e.g. risk preferences) our specific model for inertia does not matter for that analysis. The baseline model, without inertia, does yield quite different preference estimates, suggesting the inclusion of inertia is important to get correct estimates of baseline preferences.

detailed discussion of identification and estimation after we discuss all the models in this section.

Finally, our model of inertia assumes that consumers are myopic and do not make dynamic decisions whereby current choices would take into account inertia in future periods. There are several arguments to support this approach. First, price changes are not signaled in advance and, during the several years we study there are no significant premium changes. Second, it is unlikely that most consumers can forecast substantial changes to their health statuses more than one year in advance (or that they would base insurance plan choices now on these long run projections). If beliefs about long-run prices and health status are similar to present, even if some dynamic considerations exist they should not impact preference estimates markedly.

Full Model: Information Frictions. The baseline model and baseline model with inertia resemble the types of frameworks examined previously in the structural literature on health insurance markets. Our full model builds on this work by adding richer structures for consumer information (about plan features and health expenditure risk) and perceived plan time/hassle costs. While the insight that these factors matter for consumer choice is not new, the ability to cleanly measure them and disentangle them from risk preferences and health risk empirically is the central innovation of this paper. This is possible in our setting because of the rich individually-linked survey, claims, and choice data.

Our full model builds on our baseline model with the key difference being the introduction of information frictions and time/hassle costs. We introduce this additional factors by modifying the value m_j of being in plan j. m_j is impacted uniformly across all potential health state realizations from F, implying that information frictions and expected time/hassle costs impact the fixed up front utility for each plan (similar to a shift in the premium). This means that these factors impact choices at the time of plan choice but do not represent state-contingent utility factors ex post. In the full model m_j is:

$$x_{kjt} = W_{kt} - P_{kjt} - s + \eta(X_k^B)\mathbf{1}_{j_t=j_{t-1}} + \Sigma_{s=1}^{S}\beta_s I_s * I_{HDHP} + \epsilon_{kjt}$$

Here, $\mathbf{I}_s$ is an indicator variable for a given information friction or hassle cost measure, corresponding to a distinct survey question answer. S reflects the total number of such effects included

in the model. $\mathbf{I}_{HDHP}$ is an indicator variable taking on value of one if plan j is the $HDHP$ and β_s represents the incremental value the consumer places on the $HDHP$ if they have a specific information friction or expectation for time/hassle costs.

In the full model we include the following choice factors in the set of effects I_S:

- **Information about plan design features:** Our first measures examine whether a person has correct information about the $HDHP$ plan financial characteristics or not. Specifically, we construct a binary variable with value 0 if a person knows the deductible, coinsurance rate, and out-of-pocket maximum for the $HDHP$ plan and a value 1 otherwise (they don't know one of these features).[13] Our second indicator variable has value equal to one when a consumer answers 'not sure' when asked about the above plan financial characteristics and zero otherwise.

- **Provider Network Knowledge:** Our next group of measures study consumer information about the providers that can be accessed in network for each of the two plans. The first variable has value 1 if the consumer believes that one can access more providers / services in the PPO. The second variable equals 1 when the individual believes that one can access more providers through the $HDHP$ plan, while the third equals 1 if the consumer answers 'not sure' to the question on relative provider access. The omitted case for this set of indicator variables (0 for all indicators) is knowing correctly that one can access the exact same network in each plan.

- **Information about own health expenditures:** Next, we include a group of indicators to measure whether a person correctly understands their own potential health expenditures. To do this, we categorize how an individual's answer about what their expenditures were in the prior year compares to their actual expenditures during that time frame. We use three indicator variables with values equal to 1 if consumers (i) overestimate (ii) underestimate or (iii) are not sure about their actual past year of health expenditures. The omitted case is correct knowledge of past expenditures. We use this measure of past expenditure knowledge to proxy for over or under estimation of projected expenditures for the coming year (the relevant choice object).

[13] As discussed in the section on descriptive evidence, the survey questions are multiple choice questions, not recall questions.

- **Tax Benefits Knowledge:** We include a variable to measure whether or not a consumer understands the tax savings and benefits that a Health Savings Account provides (it's main advantage). Our survey asks a simple question designed to understand whether or not a consumer understands the savings that an HSA provides from foregone taxes if contributions are spent on health care expenditures. The first indicator equals one if the person answers the multiple choice question incorrectly, while the second one equals one if the person answers 'not sure.' The omitted case is the one where the person does understand the tax benefits of the HSA.

- **Time and Hassle Costs:** Our final set of measures focuses on expected time and hassle costs and the preferences that consumers have for avoiding them. Time and hassle costs differ from information frictions in the welfare analysis because they could constitute a welfare relevant choice factor if stated expected time and hassle costs equal the true expectation of time and hassle costs. Moreover, our survey measures both preferences for time and hassle costs as well as the expected quantity of them for the $HDHP$. Given the rich data, I_s for time and hassle costs equals the product of indicators for different time/hassle cost tastes and the stated expected quantities. Thus, a person who expects to incur 10 hours of time/hassle costs and answers that they have strong aversion to those costs has an indicator representing the strong aversion multiplied by ten. β_s thus measures the incremental utility from the $HDHP$ plan for each additional hour of expected time/hassle costs *conditional on stated preferences for them*. These slopes are measured relative to a baseline slope where someone answers that they have a weak preference for avoiding time and hassle costs.

Given these measures for information frictions and time and hassle costs, we estimate a series of models. First we estimate the inertial baseline model plus one additional choice factor for each of the five factors above, implying five such 'incremental' models. Then, we estimate our 'full' model where we include all five choice factors as specified above on top of the inertial baseline model. The estimates from these models can then be compared to those from both the baseline model and the inertial baseline model to see how risk preference estimates are impact by the inclusion of these additional factors. In addition, we can use these models to assess how each information friction and time/hassle costs impact consumer welfare through their impacts on plan choices.

It is important to note that the empirical model for information frictions and time/hassle costs incorporates these factors in a non-structural manner since the survey answers link directly to a change in the monetary valuation for the $HDHP$ relative to the PPO. Alternative approaches could treat these factors more structurally: for example as in the model for information frictions in section 2 we could model uncertainty about the deductible as an additional component of structural uncertainty in the expected utility model. Consumers' biased perceptions of total medical expenditures could be incorporated directly into the cost model. We feel that the reduced form representation of these factors in the model presented here captures the implications for these factors on relative plan valuations without adding the many additional assumptions that would be necessary to fully structurally model these components. The model parsimoniously captures the impact of these factors on risk preferences, conditional on baseline health risk, and the coefficients on the information frictions represent their ultimate impact on plan valuation, regardless of the true underlying model. This allows us to differentiate these factors from risk preferences in welfare analysis, our ultimate objective, without an unnecessarily complicated empirical setup.

There are several important additional modeling decisions to discuss. First, we have not modeled correlations between information frictions and either risk preferences or inertia. In addition, the model does not model correlations between health risk and risk preferences. Prior work (e.g. Cohen and Einav (2007) and Einav et al. (2010b)) has illustrated that correlations between risk preferences and and health risk can be important to model, especially when thinking about questions related to adverse selection. In our setting, unobserved heterogeneity in inertia could be related to survey question answers: if someone is uninformed about the HSP plan it could be because of inertia or unrelated information frictions. Here, our model for inertia accounts for mean impact of inertia based on observed demographics, but not unobserved heterogeneity. Future specifications will take this relationship explicitly into account. Finally, our model setup explicitly assumes that information frictions are not correlated with unobserved heterogeneity in risk preferences. Our future work will explore additional ways of categorizing consumers into information friction and hassle cost 'types' that will synthesize the different factors included here into a more parsimonious representation that takes into account strong correlation patterns in answers to survey questions. These advances will have a direct impact on the ease of incorporating some of the correlations

discussed here.[14][15]

Cost Model. The empirical choice framework, for all the specifications presented, takes the distribution of future out-of-pocket expenditures for each family, $F_{kjt}(\cdot)$, as given. This section summarizes the empirical model we use to estimate $F_{kjt}(\cdot)$. Online Appendix A presents a more formal description of the model, its estimation algorithm, and its 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), Handel (2012), 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.

The model is set up as follows:

1. For each individual and open enrollment period, we use the past year of diagnoses (ICD-9), drugs (NDC), and expenses, along with age and gender, to predict mean total medical expenditures for the upcoming year. This prediction leverages the Johns Hopkins ACG software package and incorporates medically relevant metrics such as type and duration of specific conditions, as well as co-morbidities.[16] We do this for four distinct types of expenditures: (i)

[14] One additional assumption we make is that people have similar information about the PPO option and their answers to survey questions about the HSP plan represent the relative difference in information about the HSP and PPO. This could be thought of as assuming the everyone has full information about the PPO plan, which is likely reasonable since the plan design is extremely simple and the plan has been in place for many years.

[15] We also represent the 'not sure' answer to information frictions in a reduced form way here. While the structural interpretation of that is more difficult than a wrong or right answer, if 'not sure' represents an additional level of uncertainty or disinterest above and beyond those right or wrong answers than including it separately, as we do, is relevant for plan choice predictions.

[16] 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 a time-limited acute condition, such as a broken arm.

hospital/outpatient (ii) pharmacy (iii) mental health and (iv) pharmacy.

2. We group individuals into cells based on mean predicted future utilization. For each expenditure type and risk cell, we estimate a spending distribution for the upcoming year based on ex post observed cost realizations (this takes into account within-cell variation in demographics). We combine the marginal distributions across expenditure categories into joint distributions using empirical correlations and copula methods.

3. We reconstruct the detailed plan-specific mappings from total medical expenditures to plan out-of-pocket costs.[17] We combine individual total expense projections into the family out-of-pocket expense projections used in the choice model, F_{kjt}, taking into account family-level plan characteristics.

The cost model assumes that there is no private information and no moral hazard (total expenditures do not vary with j). 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 consumers' total relative valuations of the two plans. For private information, we should be less concerned than most 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.[18]For moral hazard, Chandra et al. (2010) presents a recent review of the experimental and quasi-experimental literature, where the price elasticity for medical care generally falls in the range -0.1 to -0.4. Recent work by Einav et al. (2011), with data similar to that used 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

[17]For the two plans we study here, the different types of medical expenditures are treated similarly from a financial perspective. Modeling the four different categories of expenditures yields more precise predictions that a model based just on total medical utilization, while this expenditure division could also be useful for examining counterfactual plan design where these categories are treated differently financially (e.g. different coinsurance).

[18]Pregnancies, genetic pre-dispositions, and non-coded disease severity are possible examples of private information that could still exist. Cardon and Hendel (2001) find no evidence of selection based on private information with coarser data while Carlin and Town (2009) use similarly detailed claims data and also argue that significant residual selection is unlikely. Importantly, it is also possible that individuals know *less* about their risk profile than we do, which we account for in our full model.

hazard impact (i) is small relative to the overall difference in consumers' plan valuations and (ii) does not markedly impact our parameter estimates.

Identification. Given the individual-level claims data, choice data, and survey data, identification of the empirical parameters is relatively straightforward. In the baseline model, inertia, information frictions, and time and hassle costs are assumed away. The cost model generates a health expenditure risk distribution for each family, and, subject to that distribution, a family's choice in each year identifies a range of feasible risk preferences. With infinite variation in prices or health risk, each family's risk preference coefficient is non parametrically identified.[19] In our estimation, described in the next section, we assume a parametric form for the population distribution of risk preferences (conditional on demographics), which leads to point identification of this distribution. This is the same as identification of risk preferences throughout the literature (see e.g. Cohen and Einav (2007), Einav et al. (2011), or Einav et al. (2010a) for a survey).

Our baseline model with inertia separately identifies risk preferences from inertia by leaning on the comparison between new employees, who must always make 'active' plan choices when they arrive, and existing employees who have a default option of their previously chosen plan if they take no action. Given this, our model assumes that new employees have no inertia, while existing employees do. Table 4 describes the sample of 2339 new employees for 2011 and repeats statistics for the full population from column 1 of table 1 for comparison. New employees are slightly more likely to choose the $HDHP$, likely to be younger, likely to be lower income, and more likely to be single. Importantly, new employees span the ranges of age, gender, and income seen in the full population with non-negligible mass, such that estimates of preferences based on observable heterogeneity can credibly be extrapolated from one group to the other.

Together with the assumption that the distribution of risk preferences in the population is identical for two groups with the same demographics, the distinction between new and existing employees identifies inertia separately from risk preferences. We discuss our implementation in more detail in the estimation section: intuitively, choices from active employees identifies risk preferences conditional on demographics, and the differences between their choices and those of

[19] This is true families have the same risk preferences over time. The ϵ rationalizes choices that are 'inconsistent' with one underlying CARA value over time. Without this, the CARA model could be rejected by the data.

New vs. Existing Employees	Existing Employees	New employees
No. Employees	41,361	2339
2011 PPO%	88.8	85.7
Gender (% Male)	76.4	77.5
Age		
18-29	8.6%	36.7%
30-39	41.1%	36.3%
40-49	38.1%	20.4%
50-59	10.9%	6.2%
$\geq$60	1.3%	0.5%
Income		
Tier 1 (< \$75K)	2.7%	7.1%
Tier 2 (\$75K-\$100K)	10.1%	28.1%
Tier 3 (\$100K-\$125K)	35.3%	36.3%
Tier 4 (\$125K-\$150K)	30.5%	20.8%
Tier 5 (\$150K-\$175K)	12.0%	5.5%
Tier 6 (\$175K-\$200K)	4.7%	1.3%
Tier 7 (\$200K-\$225K)	2.0%	0.4%
Tier 8 (\$225K-\$250K)	0.7%	0.3%
Tier 9 (> \$250K)	0.8%	0.2%
Family Size		
1	23.0%	44.0%
2	19.0%	17.8%
3+	58.0%	38.2%

Table 3: This table compares employees who are new to the firm 2011 to those present in 2011 who joined the firm prior to 2011. The distinction between new employees and existing employees is central to the identification of inertia in the models described in section 4.

similar looking existing consumers identifies inertia.[20]

Building on the identification of these models, identification of the model with information frictions and hassle costs is straightforward. We observe survey data at the individual-level which we use to construct measures of these additional factors. These measures enter as observable information to the econometrician (as described in the previous section), which we assume is uncorrelated with unobserved heterogeneity in inertia or underlying risk preferences (see the discussion in the previous section for a more in depth discussion of these assumptions). Given this setup, observables differences in information frictions and hassle costs link directly to the monetary value of the plan choices actually made, conditional on risk preferences and health risk, which depend on observed demographics and health claims. The survey data thus separately identifies information frictions and hassle costs from risk preferences and inertia, while risk preferences and inertia are identified from one another as discussed above.

Estimation. In the primary implementation for each model we assume that the random coefficient γ_k for risk preferences is normally distributed with a mean that is linearly related to observable characteristics X_k^A:[21]

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

In the primary specifications X_k^A contains employee age, gender, and income.[22] We also investigate a robustness check with the assumption that γ is log-normally distributed to ensure that our main results are not specifically due to the normality assumption for γ. Finally, for all primary specifications we assume that family-plan-time specific error terms ϵ_{kjt} are i.i.d normal for each j with zero mean and variance σ_{ϵ_j}. We normalize the value of ϵ_{PPO}, the preference shock for the PPO plan, to zero and estimate the preference shock variance the $HDHP$ plan relative to that of

[20] Handel (2012) identifies inertia separately from unobserved heterogeneity in risk preferences by observing an entire population that makes an active choice in one year, and then that same population making choices in an inertial setting in subsequent years, with the plan designs and health risks changing over time.

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

[22] While age and income do change over the two years of choices the model is estimated on, they vary minimally so we treat these quantities as fixed over time and base them on the values at the beginning of 2012.

the PPO plan. [23][24]

For both the inertial baseline model and the models with inertia, information frictions, and hassle costs we assume that inertia, $\eta(X^B_{kt}, Y_k)$, is related linearly to Y_k and linked choices and demographics X^B_{kt}:

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

X^B_{kt} includes income, age, gender, and family insurance coverage tier dummies (related to number of dependents covered).

We use two different samples to estimate the sequence of models described. The first is the 'full population' sample described in column 1 of table 1, which corresponds to all employees and dependents at the firm. The second is the 'survey resampled' population described in column 2 of table 1, which corresponds to the sample of survey respondents re-weighted to fit full population plan choices.[25] Both the 'full population' sample and the 'survey resampled' sample can be used to estimate the baseline model and the baseline model with inertia. The models with information frictions and hassle costs can only be estimated with the survey resampled population, because our measures of these factors come directly from the survey question answers.

For the inertial baseline model, we use the population of new employees to estimate risk preferences and the error shock for a group of employees who are making active choices, and then use these estimates as inputs into the inertial baseline model with existing employees. Specifically, we estimate $(\mu, \delta, \sigma^2_\gamma, \sigma^2_\epsilon)$ for new employees in the baseline model and input these estimates into the inertial baseline model for existing employees as $(\widehat{\mu}, \widehat{\delta}, \widehat{\sigma^2_\gamma}, \widehat{\sigma^2_\epsilon})$ such that:

$$u_k(x) = -\frac{1}{\widehat{\gamma_k(X^A_k)}} e^{-\widehat{\gamma(X^A_k)}x}$$
$$x_{kjt} = W_{kt} - P_{kjt} - s + \eta(X^B_k)\mathbf{1}_{j_t = j_{t-1}} + \widehat{\epsilon_{kjt}}$$

[23]Since the model is a 'realized' utility model in dollar units, we don't need an ϵ variance scale normalization, just the mean normalization for the PPO.

[24]Since the PPO and $HDHP$ are only separated on financial dimensions, and not horizontal dimensions such as provider networks, we also follow Einav et al. (2011) and study a robustness specification without ϵ_{kjt}.

[25]Recall, as described in section 3, the survey was sent out to a group where $HDHP$ enrollees were over-weighted to ensure that we observe a large enough sample size of $HDHP$ enrollees in the survey.

In practice, the 'survey resampled' population is not well suited to estimating the models with inertia because only a small percentage of the overall population are new enrollees each year (approximately 5%). As a result, the full population has over 2,000 new enrolls for 2011 while the survey resampled group has less than 100. Given this, we proceed by estimating the inertial baseline model with the 'full population' and using the resulting inertia estimates as an input into the 'survey resampled' inertial baseline model and information frictions / hassle cost models. Specifically, we estimate $\eta(X_k^B)$ with the 'full population' sample, and input those estimates as $\widehat{\eta(X_k^B)}$ into the models for the 'survey resampled' population that use inertia. Thus, the state-specific money at stake in the model with information frictions and hassle costs is:

$$x_{kjt} = W_{kt} - P_{kjt} - s + \widehat{\eta(X_k^B)}\mathbf{1}_{j_t=j_{t-1}} + \Sigma_{s=1}^{S}\beta_s I_s * I_{HDHP} + \epsilon_{kjt}$$

Note that in the survey resampled models $(\mu, \delta, \sigma_\gamma^2, \sigma_\epsilon^2)$ are all estimated for that sample and each model: only inertia is input from the inertial baseline model because the 'full population' has a large population of new employees making active choices, which is important to identification of inertia.

All specifications are estimated with a random coefficients simulated maximum likelihood approach similar to that summarized in Train (2009). This approach simulates many values for the random coefficient γ and ϵ given proposed parameters for those distributions, and searches for the parameters that optimize the fit between the choices predicted by the models and the actual choices made. No simulation is necessary for coefficients related to inertia, which is estimated based on observable heterogeneity, and the coefficients for information frictions and hassle costs, which are linked directly to the observable survey questions answers. Since the estimation algorithm is similar to a standard approach, we describe the remainder of the details in Online Appendix C.[26]

5 Results

Table 5 presents the results for the inertia baseline model estimated with the full population. We present this first because it is the only model estimated with the full population sample. Subsequently, we discuss all the models estimated with the survey resampled population to present

[26]The estimation algorithm used here is similar to that used in many empirical literatures: Handel (2012) uses a similar approach in the empirical insurance literature.

a clear comparison between models to assess the impact of including information frictions and hassle costs on estimated risk preferences.

Column 1 presents the estimates for the baseline model estimated for new employees, while column 2 presents the estimates of inertia for existing employees. Recall from the estimation section that we implement the inertial model in the full population by estimating preferences for new employees in an active choice setting and use those as an input into a choice model for existing employees, using specifically to estimate inertia subject to those preferences. New employees make choices that suggest substantial risk aversion though there is a good deal of heterogeneity in risk preferences. We defer further discussion of risk preference estimates until we get to the models with the survey resampled population, so comparisons between models is an 'apples to apples' comparison. The most important figures are the estimates of inertia since these are used as an input into the survey resampled models with inertia.[27] The average amount of money foregone in plan choice due to inertia \$1,847. The impact of inertia becomes stronger with age and income while female employees have slightly lower inertia than males. Inertia has a larger impact the more dependents an employee covers. Figure 5 presents a histogram showing the distribution of estimated inertia in the population as a function of observable heterogeneity.

Table 5 presents the results for the baseline model and the baseline model with inertia for the survey resampled population. In addition to listing the estimated CARA risk preference parameters, the table also provides a simpler interpretation for those parameters for expositional purposes. The row labeled 'Gamble Interpretation of Average μ_γ' presents the value X that makes a consumer indifferent between the status quo (accepting no gamble) and accepting a gamble where they win $\$1,000$ with 50% chance and lose $\$X$ with 50% chance. Thus, if $X = 1,000$ then a consumer is risk neutral, and if $X = 0$ a consumer is infinitely risk averse. For the baseline model with no inertia X for the average consumer equals 223.04, implying substantial risk aversion, while for the baseline model with inertia the average X is 662.22. Thus, omitting inertia from the model makes consumers seem much more risk averse than they actually are, which has direct implications for counterfactual and welfare analyses.

In both models, risk aversion is negatively correlated with employee age. In the model with

[27] As discussed in the estimated section, it is important to estimate inertia with the full population, though the rest of our models are estimated with the survey resampled population, because the full population contains many new employees (> 2000) while the survey respondent population contains less than 100.

Inertia Baseline Model **Full Population**		
Sample	(1) New Employees w/o Default Opt.	(2) Existing Employees w/ Default Opt.
μ_γ - Intercept	$1.55 \cdot 10^{-2}$	-
μ_γ - Slope, Age	$-2.36 \cdot 10^{-5}$	-
μ_γ - Slope, Female	$-2.59 \cdot 10^{-4}$	-
μ_γ - Slope, Income	$-9.46 \cdot 10^{-4}$	-
Average μ_γ	$1.06 \cdot 10^{-2}$	-
Gamble Interpretation of Average μ_γ	65.39	-
σ_γ	$3.19 \cdot 10^{-2}$	-
σ_ϵ, HDHP	847.82	-
Inertia - Intercept	-	-71.47
Inertia - Slope, Age	-	28.79
Inertia - Slope, Female	-	-14.18
Inertia - Slope, Income	-	124.26
Inertia - Slope, Family size = 2	-	335.28
Inertia - Slope, Family size > 2	-	370.79
Average Inertia	-	1847.58

Table 4: This table presents the results from the first-stage of our estimation. The model estimates preferences for new employees, making active plan choices, in 2011, and estimates inertia by comparing the choices of existing employees to those of new employees subject to these preference estimates. The inertia estimates here are then used as an input into all models estimating inertia with the the survey re-sampled population.

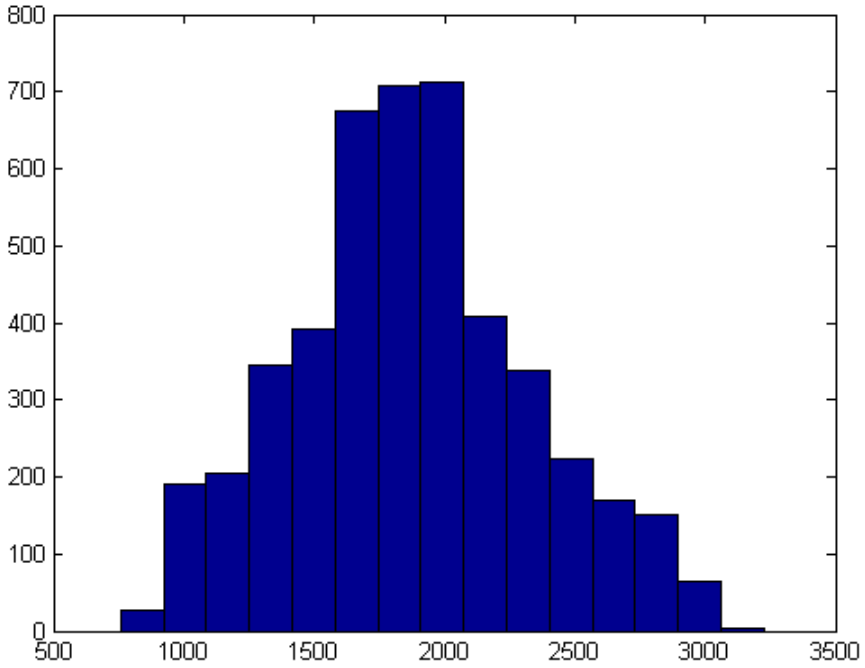


Figure 10: Histrogram of Inertia Estimates with Observable Heterogeneity

inertia it is negatively correlated with being female and negatively correlated with income, though both of these latter effects are relatively small. In addition to this observable heterogeneity, there is substantial unobservable heterogeneity in both models: in the model with inertia σ_γ is about the same magnitude as the average μ_γ.[28] For the model with inertia $\sigma_\epsilon = 236$ so much variation in choices is explained without the idiosyncratic error shock in that specification.

Table 5 presents the results from our main models, which contain information frictions plus time and hassle costs. Column 1 repeats Column 2 from table 5 with the results from the inertial baseline model. Columns 2-5 present results from the incremental information friction models where we add only one set of survey answers related to one information friction or to hassle costs. Column 6 presents the results of the full model with all information friction measures and hassle cost quantities and tastes. To interpret the CARA coefficient estimates we use the same gamble interpretation discussed above for table 5. Compared to mean $X = 662$ for the inertial baseline model, mean $X = 707$ for the model that includes information frictions related to plan design features and 711 for the model with information frictions related to own total medical expenditures. Thus, in our

[28]Recall, since the risk preference random coefficient is assumed normal distributed, the simulated values of γ are truncated just above zero, such that in our setting there is a mass of risk neutral consumers.

Baseline Models **Survey Respondents Resampled**	(1) No Inertia	(2) Inertia
μ_γ - Intercept	$6.88 \cdot 10^{-3}$	$1.00 \cdot 10^{-3}$
μ_γ - Slope, Age	$-1.07 \cdot 10^{-4}$	$-1.11 \cdot 10^{-5}$
μ_γ - Slope, Female	$1.98 \cdot 10^{-4}$	$-3.87 \cdot 10^{-5}$
μ_γ - Slope, Income	$683 \cdot 10^{-5}$	$-1.18 \cdot 10^{-5}$
Average μ_γ	$2.99 \cdot 10^{-3}$	$5.03 \cdot 10^{-4}$
Gamble Interpretation of Average μ_γ	223.04	662.22
σ_γ	$5.20 \cdot 10^{-3}$	$6.75 \cdot 10^{-4}$
σ_ϵ, HDHP	502.23	236.68

Table 5: This tables presents the structural estimates from the baseline model and the inertial baseline model. The former models health risk, risk preferences, and idiosyncratic preferences while the latter models those factors plus inertia, which is an input simulated from the first-stage estimation presented in table 5.

setting when we include these additional choice relevant factors consumers are estimated to be less risk averse, so a model that omits these factors over-predicts risk aversion here. The model that incorporates information frictions related to provider network knowledge has mean $X = 597$ on so predicts slightly higher risk aversion than the model without those frictions. Finally, the incremental model that includes measures of time and hassle costs (which are potentially welfare relevant) predicts mean $X = 817$ so omitting measures of time and hassle costs has a significant effect on risk preference estimates. The full model with all additional choice frictions has mean $X = 853$ so predicts the lowest extent of risk aversion. Thus, in our setting having the linked survey data to proxy for information frictions and hassle costs has a substantial impact on estimated risk preferences. We discuss the implications for welfare analysis and counterfactual analyses in the next section.

For the information frictions themselves, in the full model having any plan design question incorrect about the $HDHP$ causes the consumer to value the $HDHP$ by \$183 less than a consumer

with full information. The incremental model predicts a value of -$234 for this factor. In the full model, consumers who believe that the PPO plan has a larger network of medical providers value the $HDHP$ by -$1,787 less than someone who knows that these plans have the exact same provider network (for the incremental model this same coefficient is -$644). Those who underestimate their own total medical expenditures for the past year value the $HDHP$ by $425 less than someone who knows their recent expenditures, reflecting the fact that generally uninformed consumers might be hesitant to choose the $HDHP$. This is further reflected by the fact that consumers answer they are 'not sure' about their past medical expenditures are value the $HDHP$ by $262 less than those with correct information.[29] The incremental models for total medical expenditure knowledge predict negative values slightly lower in magnitude for these two coefficients.

Of all the additional frictions, stated time and hassle cost quantities and preferences seem to have the most substantial impact on choices. For each additional stated hour of time spent on plan billing, administration, and logistics, in the full model a consumer values the $HDHP$ by $120 less if they 'accept but are concerned' about hassle costs. Those who state they have a strong distaste for hassle costs value the $HDHP$ plan by $117 less for each additional stated hour of these tasks. Reassuringly, those who state that they are not particularly concerned about hassle costs do not value the $HDHP$ plan by less for each incremental stated expected hour spent (the point estimate per hour suggests you prefer the $HDHP$ by $19 more for each hour). The values for all of these coefficients is similar in the incremental model with only time and hassle cost measures added to the inertial baseline model.

Overall, the results reveal that our measures of information frictions and hassle costs matter for plan choices, in some cases leading consumers to value the $HDHP$ by more than $1000 less than individuals without those given frictions. We believe these results illustrate (i) that these frictions enhance choice predictions relative to models without them and (ii) shed light on which factors may be most important to consumers making choices in complex environments with uncertainty. Perhaps most importantly, the results illustrate that *not including* these frictions can substantially bias risk preference estimates, which in turn has implications for welfare analyses and counterfactual investigations. Our next section highlights this point in the context of a menu design counterfactual,

[29] In future work we will study correlations across the information frictions and explore grouping consumers into information 'types.' Our current results suggest that consumers who are uninformed in any way tend to prefer the easy to understand PPO plan than the harder to understand $HDHP$ plan.

Information Frictions & Hassle Costs Survey Resampled						
Add. Frictions	(1) None	(2) Plan Design Knowledge	(3) Time / Hassle Costs	(4) Provider Networks	(5) TME Info	(6) Full Model
Average μ_γ	$5.03 \cdot 10^{-4}$	$4.10 \cdot 10^{-4}$	$2.23 \cdot 10^{-4}$	$6.61 \cdot 10^{-4}$	$4.04 \cdot 10^{-4}$	$1.71 \cdot 10^{-4}$
Gamble Interpretation of Average μ_γ	662.22	707.08	817.12	596.93	710.50	853.94
σ_γ	$6.75 \cdot 10^{-4}$	$5.25 \cdot 10^{-4}$	$2.26 \cdot 10^{-4}$	$8.95 \cdot 10^{-4}$	$5.14 \cdot 10^{-4}$	$1.52 \cdot 10^{-4}$
σ_ϵ, HDHP	236.68	20.25	119.52	609.00	126.12	542.19
Benefits knowledge:						
Any incorrect	-	-233.79	-	-	-	-182.88
Any 'not sure'	-	41.98	-	-	-	-321.36
Time cost hrs. X prefs:						
Time cost hrs.	-	-	20.41	-	-	19.75
... X Accept, concerned	-	-	-88.63	-	-	-120.98
... X Dislike	-	-	-103.09	-	-	-117.72
Provider networks:						
HSP network bigger	-	-	-	-145.27	-	-394.54
PPO network bigger	-	-	-	-644.02	-	-1787.64
Not sure	-	-	-	240.85	-	-148.82
TME guess:						
Overestimate	-	-	-	-	-210.48	-53.59
Underestimate	-	-	-	-	-195.16	-425.78
Not sure	-	-	-	-	-102.64	-262.90
Tax benefits:						
Incorrect	-	-	-	-	-	288.55
Not sure	-	-	-	-	-	613.78

Table 6: This table presents the results from the primary models with information frictions and hassle costs. Column 1 repeats the inertial baseline model results from column 2 of table 5, which models health risk, inertia, risk preferences, and idiosyncratic plan tastes. The incremental models presented in Columns 2-5 add either a specific information friction or time/hassle costs to the inertial baseline model, as described in section 4. Column 6 presents the results of our full model, which includes all information frictions and time/hassle costs.

an oft-studied aspect of insurance markets in the literature.

6 Policy Analysis: Welfare Impact of Forced HDHP Switch

We assess the welfare implications of incorporating information frictions and hassle costs into the model of consumer choice by investigating the welfare impact of a health insurance plan menu redesign. We focus on the case where the firm completely removes the the PPO option from the choice set and forces consumers to enroll in the $HDHP$ plan. This is an interesting analysis because for 2013, the firm we study decided to change the menu in this exact way, forcing over 40,000 employees to change from the PPO to the $HDHP$. Outside the context of the large firm we study, other employers face similar menu design choices, and many have made similar choices to move their employees towards high-deductible plans in recent years. From a public policy perspective, health insurance exchange regulators must decide on which actuarial equivalence values (degree of consumer cost sharing) to allow private insurance companies to offer to consumers. Our analysis illustrates how welfare analyses for these analysis will be changed by incorporating information frictions and hassle costs, which our previous section reveals can have substantial impacts on consumer choices.

Specifically, we compute the welfare impact of our firm switching from a menu with the $HDHP$ and PPO options to one with just one choice, the $HDHP$ option. We investigate the welfare impact for three models with differing levels of detail on choice relevant factors: (i) the baseline model (ii) the inertial baseline model, and (iii) the full model with all information friction and hassle cost measures included. We use the survey resampled population as the basis for this analysis since this is the sample all three models can be estimated for. Since the analysis studies movement from a menu offering to a forced choice, including inertia in the model is only relevant insofar as it impacts estimates of risk preferences and other factors that are relevant to studying welfare of actually being enrolled in the $HDHP$.

We conduct the welfare analysis for these three models to reveal exactly how the welfare implications of this simple menu redesign change when we are able to measure information frictions and hassle costs in greater detail. In the baseline model and inertial baseline models, which are similar to those used elsewhere in the literature, risk preference and health risk estimates are the

primary drivers of welfare analysis. In our richer model with information frictions and hassle costs, these factors impact choices in a specific way, and change the risk preference estimates relative to the baseline model where these factors were omitted. This has direct and potentially significant implications for welfare analysis: if information frictions are the primary reason consumers choose a plan instead of risk preferences this matters a lot for computing the welfare impact of a forced switch to the $HDHP$ plan.

Welfare Analysis. 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_{kjt}. [30] Formally, for the baseline model and inertial baseline model Q_{kjt} is determined for each family, plan, and time period by solving:

$$= U_{kjt}(\gamma_k, \epsilon_{kjt}, F_{kjt}(\cdot)) = u(Q_{kjt}) = -\frac{1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)(W - Q_{kjt})}$$

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. In our setting, since Q_{kjt} is a certainty equivalent loss, lower values of Q_{kjt} are better from the consumer perspective. For example, $Q = 0$ implies full insurance with no premium, similar to the PPO option in our setting. Since the $HDHP$ plan gives money to consumers up front through the firm health savings account contribution, if $Q_{kjt} > 0$ then the consumer is worse off in the $HDHP$ plan, but if $Q < 0$ their certainty equivalent loss is negative and they are better off in the $HDHP$ plan.

For a general policy Z impacting health plan design and subsequent consumer health plan enrollment, the change in mean consumers surplus from the policy Z is:

$$\Delta CS_t^Z = \frac{\Sigma_K [Q_{kjt} - Q_{kjZt}^Z]}{K}$$

Here, Q_{kjZt}^Z is the certainty equivalent loss for the consumer of enrolling in the plan they actually

[30] Though our welfare analysis for the forced switch to the $HDHP$ plan only requires computing the welfare for that plan, we present a more general version of welfare analysis for enrollment in any plan to illustrate how the methodology can be used to answer other questions.

choose or enroll in in policy environment Z, given potentially new characteristics of health plans in that environment. Moving from this general notation to our specific question, the mean welfare impact of a policy that forces all consumers into the $HDHP$ is:

$$\Delta CS_t = \frac{\Sigma_K[Q_{kjt} - Q_{k,HDHP,t}]}{K}$$

Since the $HDHP$ plan design is the same before and after the forced switch, we drop the Z notation here and note that the change is just the welfare impact of switching all those formerly in the PPO to the $HDHP$, holding plan design constant.

The above model pertains to the baseline model and inertial baseline model in our context (since inertia is assumed away, even though risk preference estimates change the welfare model is the same). The model with information frictions and hassle costs adds complexity to this framework since we must model the impact of these additional choice factors on welfare. While the estimation results reveal that the information frictions can have a significant impact on choices, some are not welfare relevant when considering the question of a forced switch to the $HDHP$. For example, when deciding between the PPO and $HDHP$, someone who believes they can see more providers in the PPO is much more likely to choose that plan, even though they can actually see exactly the same providers in the $HDHP$. In this case, this information clearly matters for choice. However, when this consumer is forced into the $HDHP$ plan, the providers are actually the same so their actual welfare is not impacted upon enrollment. Welfare from actual provider access is identical across both plans, evens if information about that impacts choices.

We define I_{WS} as the set of variables representing information friction and hassle cost factors that do actually have tangible welfare implications upon forced enrollment, and $I_{\overline{W}S}$ as the set of variables that don't. I_{WS} and $I_{\overline{W}S}$ are mutually exclusive and, when combined, are equivalent to the set I_S of variables included in the full model described in section 4. For the welfare analysis, we set all variables in the set $I_{\overline{W}S}$ equal to 0: these variables, which are described below, don't factor into the welfare analysis at all so their coefficients and values are irrelevant for that analysis. Of course, they are relevant in the sense that they change the model estimates of welfare relevant factors such as risk preferences, but don't directly matter for welfare upon the forced switch.

Given this setup, the certainty equivalent of enrolling in a plan j in the full model with infor-

mation frictions and hassle costs is:

$$= \widehat{U_{kjt}}(\gamma_k, \epsilon_{kjt}, F_{kjt}(\cdot), I_{Ws}) = u(Q_{kjt}) = -\frac{1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)(W - Q_{kjt})}$$

Thus, if a choice friction does matter for welfare, it's coefficient is used as set forth in the full model in section 4 β_s as a welfare relevant factor when someone is forced to enroll in the $HDHP$. The variables $I_{\overline{W}S}$ and associated coefficients don't matter for welfare and aren't incorporated. Conditional on this new model computing Q_{kjt} in the presence of choices frictions, the mean consumer surplus calculation of the forced switched to the $HDHP$ is the same as above.

A crucial issue is to determine which frictions are included as welfare relevant given the forced switch, and which are not. In our analysis, the following choices factors are assumed to be welfare relevant:

- Risk preferences
- Health out-of-pocket expenditure risk
- Idiosyncratic preference ϵ[31]

We assume the following factors are not welfare relevant:

- Information frictions relating to health plan financial characteristics
- Information frictions relating to provider networks for both plans
- Information frictions relating to HSA tax benefits
- information frictions about own total medical expenditures
- Expected time and hassle costs and related preferences for them

We believe that, using arguments similar to that given above for why provider network information frictions, it is straightforward to argue that plan financial characteristic and own total medical expenditure information frictions should not be welfare relevant. We believe that frictions relating to HSA tax benefits fall into this category as well, though we note that if this information friction

[31] It is valid to debate whether this should be included as a welfare relevant factor or not. We plan to address this further in a future version of this paper.

causes the consumer to place less money in an HSA this friction falls somewhere in between welfare relevant and welfare irrelevant.

The most controversial of these categorizations is expected time and hassle costs in the $HDHP$ plan and related distaste for those factors. We believe that if the stated time and hassle costs of those in the PPO correspond to the time and hassle costs that those employees would actually experience once enrolled in the $HDHP$, then those stated costs *should* be counted as welfare relevant factors. However, simple comparisons among the stated expected time spent for hassle / logistical issues is a good deal higher for those in the PPO plan in our data, relative to those in the $HDHP$ plan. This suggests that either (i) there is real heterogeneity in experienced $HDHP$ hassle costs, and those who select into the PPO have higher such costs or (ii) PPO enrollees perceive that the $HDHP$ will have hassle costs that are substantially higher than those they would actually experience once forced to enroll in that plan. While in future work we intend to use features of the choice and survey data to separate actual experienced hassle costs and misperceptions about those costs for PPO enrollees, in our analysis here we include time and hassle costs as an irrelevant welfare factor. Maintaining this assumption highlights the impact that the change in estimated risk preferences in the model with choice frictions has on the welfare implications of the $HDHP$ switch. Our implementation is thus similar to assuming that consumers are forced to switch into an $HDHP$ plan that is the same as that observed in the data, but that this plan has the same experienced time and hassle costs as the PPO.

Results. The top half of table 6 presents the welfare results for the policy that removes the PPO option and forces all consumers to enroll in the $HDHP$ plan. It presents the mean and distribution of welfare implications for (i) the baseline model (ii) the inertial baseline model and (iii) the full model with information frictions and hassle costs. In addition, we present the distribution of the monetary gain for the forced switch for risk neutral consumers with no other preference factors (this can be thought of as proxying for the maximum welfare gain for switching people to the $HDHP$ given its current design).

The welfare implications of including information friction and hassle cost measures are evident when comparing the mean consumer surplus change across the models. The baseline model, where consumers are estimated to be very risk averse, predicts a mean \$1,236 consumer surplus loss

from forcing consumers into the $HDHP$. The inertial baseline model predicts a mean loss of \$862 from the forced switch, revealing that controlling for inertia is an important component of preference estimation in this setting. The full model predicts a mean \$587 loss from the forced switch, which is lower than both baseline models because including information frictions and hassle costs reduces estimated risk aversion. Crucially, as discussed above, these added frictions are, in general, important for predicting consumer choices but don't matter for welfare in a tangible way when consumers are forced to switch into that plan. Thus, even with a fairly sophisticated model with inertia and risk preferences, omitting measures of these frictions predicts a \$300 per person lower mean surplus for switching into the $HDHP$. The mean loss from the forced switch if all consumers are risk neutral and no other preferences are considered is \$497. Thus, in our setting the mean incremental welfare loss from risk aversion is approximately 3 times as large when omitting information friction and hassle cost measures from the model.

Figure 6 provides a graphical representation of the distributional welfare results contained in table 6. It presents the welfare results for the 10th, 25th, 50th, 75th, and 90th quantiles in the population for each model. The median loss from the menu redesign is \$1,465 in the baseline model, \$859 in the inertial baseline model, \$462 in the full model with all frictions, and \$301 with the simple risk neutral model. The ranking of the welfare loss remains the same for all four models across all of the quantiles examined. The relative magnitude of the differences in welfare between the models remains fairly consistent across all the quantiles examined. This distributional illustrates that (i) in all models there is substantial variation in the welfare impact of the forced $HDHP$ switch and (ii) the implications of including information frictions and hassle costs are significant regardless of whether the mean or a given quantile are used.

In order to illustrate the implications of these results from the perspective of a policymaker, or a firm's HR head, we analyze what the minimum necessary amount of moral hazard is to justify the forced shift to the $HDHP$ plan.[32]. One primary motivation for the firm to switch to the $HDHP$ is to incentivize consumers to reduce 'wasteful' medical expenditures. This motivation is a primary underlying reason for why many firms around the country are switching to high-deductible plans. However, switching to these plans does place an increased risk burden on consumers, so it is important to understand whether the social benefits from reduced wasteful expenditures outweighs

[32]Recall that while this is a counterfactual analysis, this change actually occurred at the firm for 2013

Forced $HDHP$ Enrollment Welfare Analysis					
Model	Mean	25th	Median	75th	95th
Base model no inertia	-1235.75	-2503.66	-1465.35	-174.43	362.44
Base model	-862.80	-1751.41	-859.42	0.00	713.77
Full model	-587.54	-1435.23	-462.20	0.00	1099.83
Risk neutral	-497.27	-1317.40	-301.34	121.95	1148.95

Moral Hazard Necessary To Justify Switch		
	Elasticity lower bound	Elasticity upper bound
Base model, no inertia	0.416	0.717
Base model	0.278	0.479
Full model	0.211	0.364
Risk neutral	0.157	0.271

Table 7: The top half of this table presents the welfare impact of a menu redesign that removes the PPO option and forces all consumers into the $HDHP$ plan. The welfare results are presented for each of four different models to illustrate the impact of incorporating inertia, information frictions, and hassle costs to a basic model with just health risk and risk preferences. The bottom half of the table illustrates, for each potential underlying model, the minimum consumer price elasticity of demand for medical expenditures that can generate enough cost savings to justify the forced switch to the $HDHP$ plan. Notably, the threshold for moving consumers to the $HDHP$ is quite different for the full model relative to the basic model.

the social welfare loss from increased risk exposure. Thinking about how these two dimensions trade-off should be an essential component of any regulator or HR head's calculation when thinking about social welfare.

We implement this analysis by calculating the implied savings from reduced wasteful medical expenditures across a range of potential consumer price elasticities of medical expenditures. The analysis does not take a stance on what this elasticity is, but instead is intended to find the *minimum elasticity* such that, for any elasticity above that minimum, switching everyone into the $HDHP$ plans is socially optimal. Since the different models that we study predict different mean consumer welfare losses (due in large part to different risk preference estimates when additional factors are included), they will also require different minimum elasticities to justify the full switch to the $HDHP$. As the consumer welfare loss from the forced $HDHP$ switch predicted by a given model

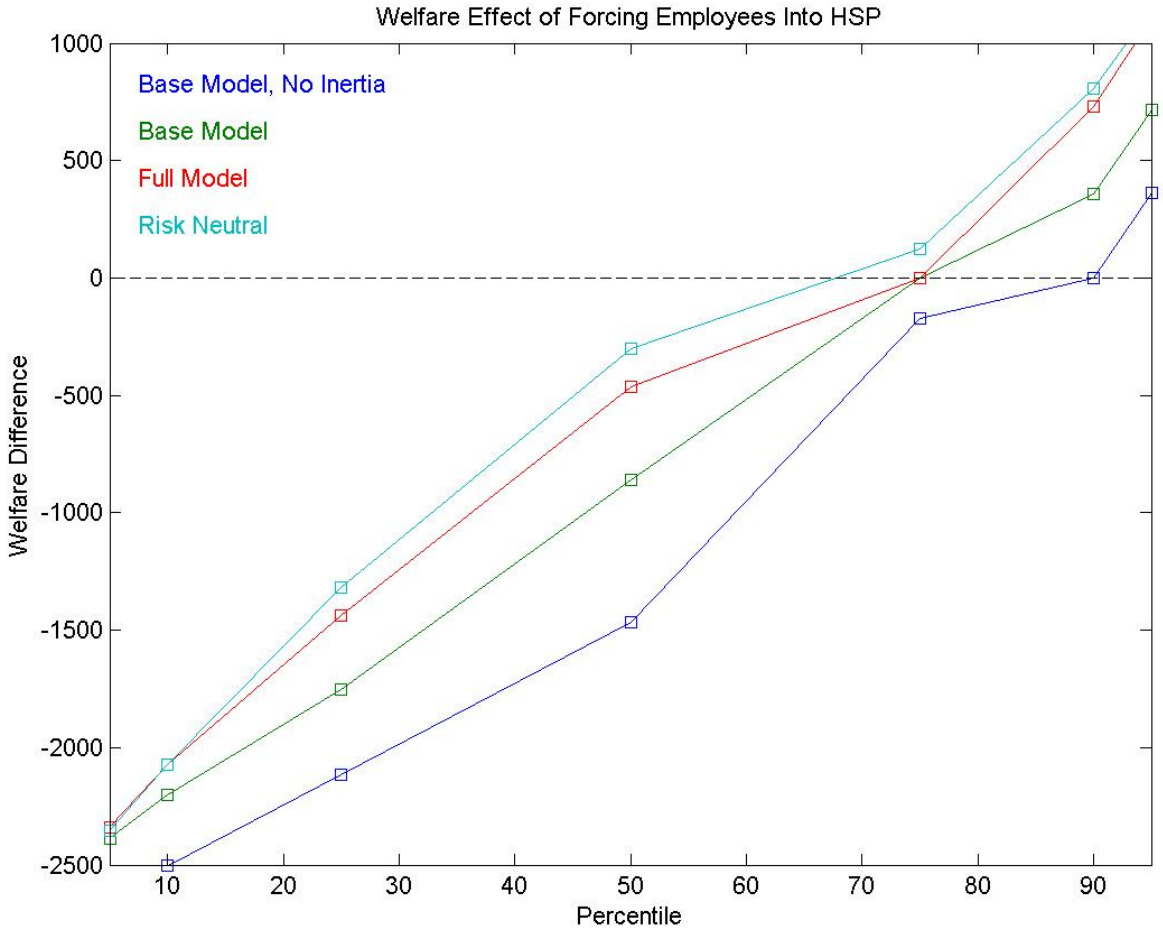


Figure 11: This figure plots quantiles of the welfare impact of the the forced $HDHP$ switch for each of the four models presented in table 6.

becomes larger, a larger price elasticity is necessary to justify that menu redesign. It is important to note that the savings from reduced total medical expenditures due to the forced $HDHP$ switch are assumed to accrue entirely to the firm: the consumer choice models estimated assume zero moral hazard so we do not explicitly model how reduced total medical expenditures reduce also reduce consumer out-of-pocket payments. To the extent that consumers' payments are reduced, we are slightly over-predicting the current minimum elasticity necessary to justify the policy change. In a future version of this work we will also specifically model how the candidate elasticity reduces consumer out-of-pocket expenditures, and incorporate that into this analysis.[33]

[33]In addition, in future work we will run a simple robustness check for the full model with information frictions that adjusts out-of-pocket spending distributions F down for the $HDHP$ plan to account for the fact that consumers

We calculate the cost savings from reduced medical expenditures in the $HDHP$ plan due to consumer price responsiveness as a function of a candidate elasticity and a derived high-level price for medical expenditures in the $HDHP$:

$$\Delta TC = \frac{\Sigma_K \mu_{F_k,HDHP}}{\Sigma_K \mu_{TME_k,HDHP}} * \frac{\Sigma_K \mu_{TME_k,HDHP}}{K} * \xi$$

Here, TC represents the total reduction in medical expenditures from forcing everyone in the PPO to switch to the $HDHP$. We note that, for this analysis, we only consider the approximately 85% of the population actually forced to switch into the $HDHP$: those who were already in the $HDHP$ are excluded because under the menu redesign they continue to enroll in the same plan and the policy does not reduce their expenditures further since they were already enrolled in the $HDHP$.[34] The first fraction on the right hand side measures the mean consumer price of medical expenditures in the $HDHP$ plan, measured as the average proportion of expenditures paid by consumers if (i) they had the same total expenditures as the PPO and (ii) they incur those expenditures under the $HDHP$ plan. Here, μ_{TME_k} denotes mean predicted total medical expenditures for family k in 2012 while $\mu_{F_k,HDHP}$ denotes mean predicted out-of-pocket expenditures for that family in that year. The second fraction determines mean total predicted medical expenditures across all families k in 2012. ξ is the assumed candidate price elasticity for medical expenditures in the population. For this analysis we keep things simple and assume a homogeneous elasticity in the population. Intuitively, the total cost savings in medical expenditures from shifting the PPO consumers to the $HDHP$ equals the marginal price difference between the PPO and $HDHP$ multiplied by the elasticity ξ (to get the proportional reduction in expenditures) and then multiplied by total mean medical expenditures to get actual cost savings.

When using these total cost savings in the context of a welfare comparison, it is also crucial to consider whether services foregone are purely wasteful or whether they have some value to consumers. If we compare ΔTC calculated above to the consumer welfare implications from the choice model, we are implicitly assuming that reduced medical expenditures come from a reduction

may have some price elasticity of demand for medical expenditures. The purpose of this analysis will be to show that even a fairly high price elasticity of demand for medical expenditures will not substantially change the main choice parameters studied or our primary conclusions.

[34] Given this, the mean welfare calculations used here exclude the 0's from people who were already enrolled in the $HDHP$ plan, who clearly experience no welfare change in the counterfactual menu redesign that leaves only that plan. These 0's are included in the overall welfare results in the top of table 6.

in purely wasteful services. In reality, if consumers utilize medical services rationally, then we know they value them more than the price they are paying. While the marginal consumer price in the PPO is always zero, if we take the marginal price in the $HDHP$ to be the average price paid, defined as $p_{HDHP} = \frac{\Sigma_K \mu_{F_k,HDHP}}{\Sigma_K \mu_{TME_k,HDHP}}$, then we know that consumers' values for services they would do in the PPO but not in the $HDHP$ lies between those two marginal prices.[35] In this simple model, we can bound the welfare loss to consumers from services foregone below $p_{HDHP} * \Delta TC$.

Given this, we can bound the minimum elasticity necessary to justify the switch to the $HDHP$ between that ξ that equates ΔTC with the change in consumer surplus ΔCS and the elasticity that equates $1-p_{HDHP}*\Delta TC$ and ΔCS. The upper bound follows from the fact that, if consumers value all services foregone just below $p_{HDHP}\Delta TC$, the the true value from reducing medical expenditures is $1-p_{HDHP}*\Delta TC$ and ΔCS. This can then be compared to consumer welfare from plan enrollment in the $HDHP$ derived from our choice model (which assumes $\xi = 0$).

The bottom half of table 6 presents the bounds on the minimum elasticity necessary to justify the forced switch to the $HDHP$ from a social welfare perspective. The first column presents the lower bound for this minimum elasticity (where all foregone spending is assumed purely wasteful) while the second column presents the upper bound on this minimum elasticity (where all foregone spending is valued just below consumers' marginal prices). Figures 6 and 6 illustrate these calculations in depth. Figure 6 plots the consumer welfare change for the forced switch to the $HDHP$ for (i) the baseline model (ii) the inertial baseline model (iii) the full model with frictions and (iv) the simple risk neutral model all as horizontal lines. The vertical axis illustrates the magnitude of the mean welfare impact predicted by those models for those switching fro the PPO to the $HDHP$. The horizontal axis represents the candidate price elasticities ξ. The diagonal line represents ΔTC as a function of ξ. As ξ increases, $\Delta TC(\xi)$ increases. When $\Delta TC(\xi)$ intersects with mean consumer welfare change predicted by a given choice model, the cost savings to the firm from switching to the $HDHP$ exceed the consumer welfare loss from switching to that plan. then, for every value above ξ, the forced switch to the $HDHP$ is welfare enhancing.

Figure 6 illustrates that the minimum elasticity necessary to justify the menu redesign under

[35]Of course, this back of the envelope calculation ignores the fact that the $HDHP$ plan is a non-linear contract where the marginal price consumers face at any point in time is unclear and depends on their expectations about their end of year spending See Aron-Dine et al. (2012) for an extensive treatment of this issue. See Baicker et al. (2012) for a discussion of cases where consumers forego care that society should value for them at higher than their marginal prices.

when foregone spending is purely wasteful equals 0.416 for the baseline model. Recall, the baseline model predicts the highest mean consumer welfare loss from switching to the $HDHP$ because it over-predicts risk preferences by omitting inertia and the additional choice friction measures. This minimum elasticity is 0.278 for the inertial baseline model, 0.211 for the full model with information frictions, and 0.157 under a simple risk neutral model. Thus, as we examine our incremental models and add inertia and choice relevant frictions, it is easier to justify the policy switching all consumers to the $HDHP$. This is because the welfare loss consumers experience is lower as biases to risk preference estimates in the baseline model are removed by including additional choice elements, and, consequently, a lower reduction in total medical expenditures caused by switching from the PPO to the $HDHP$ is necessary to justify that switch from a social welfare standpoint.

Figure 6 adds a line representing $1 - p_{HDHP} * \Delta TC$ to study the upper bound on the the minimum elasticity necessary to justify the switch to the $HDHP$ (the case where foregone services are valued just below p_{HDHP}). This line lies below ΔTC since, for a given candidate elasticity ξ the benefit to consumers reducing medical expenditures is lower. With this upper bound, the minimum elasticity necessary to justify the menu redesign is 0.717 in the baseline model, 0.479 in the inertial baseline model, 0.364 in the full model, and 0.271 for the risk neutral model. The figure includes the lower and upper bound on the minimum elasticity to justify the policy for all 4 models.

Overall, our analysis of a policy that removes the PPO option from the choice set and moves all consumers into the $HDHP$ illustrates how including measures of information frictions and hassle costs can directly impact welfare analysis and policy decisions in health insurance markets. With the fairly naive baseline model, a very high price elasticity of demand (relative to those estimated in the literature) would be necessary to justify moving all consumers to the $HDHP$. However, our framework reveals that, in our setting, this is because that model omits key choice factors such as inertia, information frictions and hassle costs, the latter two of which may impact actual choices made rather than welfare once enrolled in a plan. The full model with all of these features, which leverages our detailed survey measures for choice frictions, suggests that a much lower price elasticity in between 0.211 and 0.364 is necessary to justify the policy change. These price elasticities fall exactly in the range of those usually predicted by the literature (see e.g. Chandra et al. (2010)). Thus, while a policymaker in our setting would generally not want to switch all consumers to the

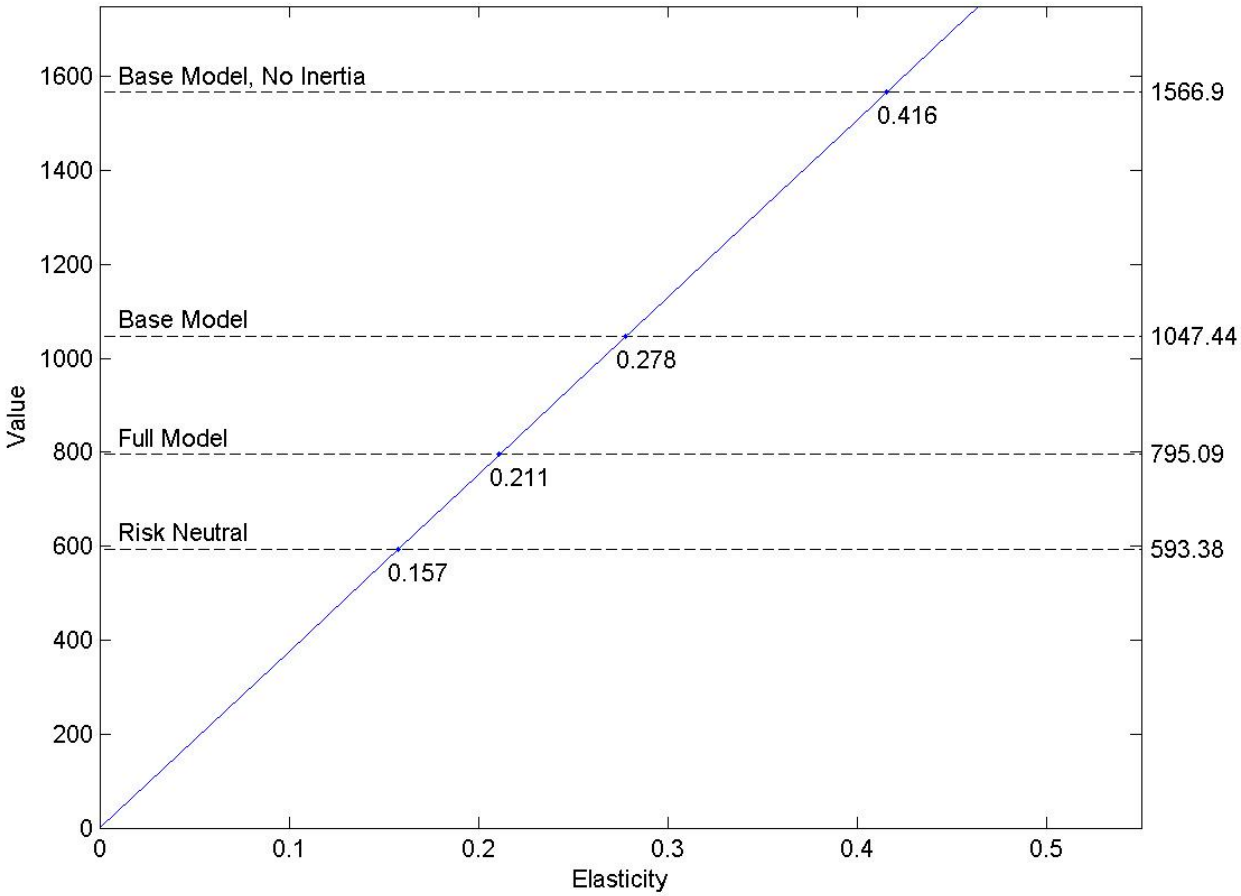


Figure 12: This figure illustrates how the minimum price elasticity necessary to justify the forced switch to the $HDHP$ is calculated for each of the models studied. The horizontal lines represent the mean consumer surplus impact of the forced switch to the $HDHP$ plan for the different models studied. The diagonal line represents the relationship between total cost savings (y-axis) from reduced expenditures in the $HDHP$ for a given candidate price elasticity of demand (x-axis). For this chart the medical expenditures foregone are assumed to be 'purely wasteful.' The place where the consumer welfare and cost savings lines intersect (for each model) determines the minimum elasticity necessary to justify the forced switch under that model.

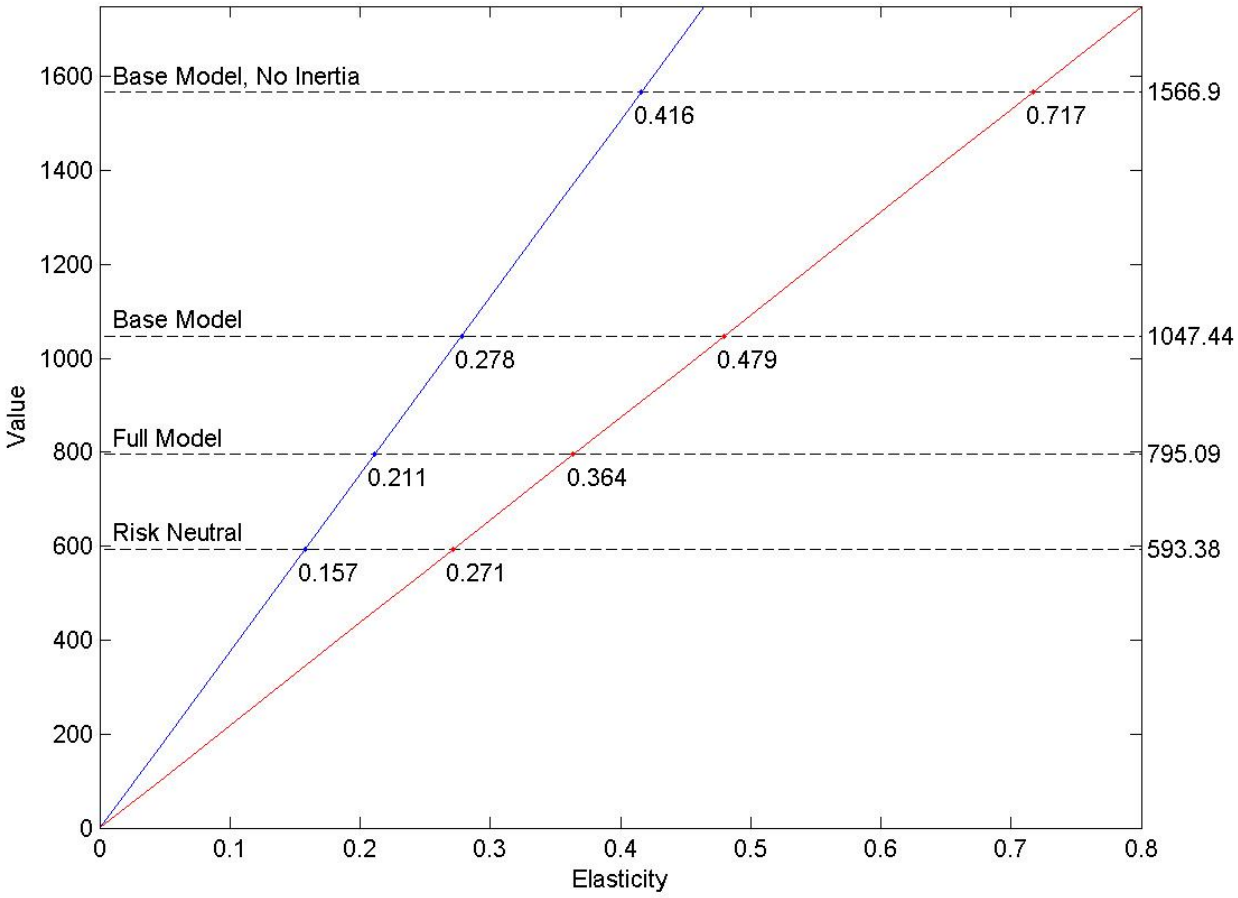


Figure 13: This figure is the same figure as 6 but also contains a second cost savings diagonal line that incorporates the fact that some foregone medical expenditures could be valuable. Under the assumption that consumers understand the value of their medical care choices, the second line represents total cost savings minus the value of foregone care if that care is valued just below the average marginal price of care paid by consumers. Thus, the two diagonal lines imply lower and upper bounds on the minimum price elasticity necessary to justify the forced $HDHP$ switch.

$HDHP$ with the baseline model, and would probably not want to in the inertial baseline model (though that choice could be marginal), in our full model with information frictions included this switch would be justifiable based on many analyses of moral hazard from the literature.

7 Conclusion

TO BE ADDED

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.

Aron-Dine, A., L. Einav, A. Finkelstein, and M. Cullen, "Moral Hazard in Health Insurance: How Important is Forward Looking Behavior?," November 2012. NBER working paper no. 17802.

Arrow, K., "Uncertainty and the Welfare Economics of Medical Care," *The American Economic Review*, 1963, *53* (5), 141–149.

Baicker, K., S. Mullainathan, and J. Schwartzstein, "Behavioral Hazard in Health Insurance," October 2012. Dartmouth Working Paper.

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.

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.

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

Cutler, D., M. McClellan, and J. Newhouse, "How Does Managed Care Do It?," *The RAND Journal of Economics*, 2000, *31* (3), 526–548.

Dube, J., G. Hitsch, P. Rossi, 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 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. NBER working paper no. 16969.

Ericson, K. and A. Starc, "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," April 2012. Boston University, Working Paper.

Handel, B., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," August 2012. University of California at Berkeley Working Paper.

Ho, K., "Insurer-Provider Networks in the Medical Care Market," *The American Economic Review*, 2009, *99* (1), 393–430.

Kling, J.R., S. Mullainathan, E. Shafir, L Vermeulen, and M.V. Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," September 2011. NBER working paper no. 17410.

Kowalski, A., "Estimating the Tradeoff Between Risk Protection and Moral Hazard with a Non-linear Budget Set Model of Health Insurance," May 2012. NBER working paper no. 18108.

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.

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

Spinnewijn, J., "Heterogeneity, Demand for Insurance and Adverse Selection," 2012. LSE.

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

A Appendix: Survey Instrument

Survey responses			
Question	Correct	Incorrect	Not Sure
What is the deductible under the HDHP?	27.08%	22.40%	50.53%
HDHP-Existing	52.68	11.23	36.10
HDHP-New	50.79	13.49	35.73
PPO	21.53	24.66	53.82
How much is the employer HDHP subsidy?	31.42	19.94	48.64
HDHP-Existing	73.40	11.05	15.54
HDHP-New	68.65	11.21	20.14
PPO	22.50	21.92	55.58
What is the out-of-pocket maximum under the HDHP?	18.47	21.98	59.55
HDHP-Existing	28.32	22.11	49.57
HDHP-New	31.87	18.91	49.21
PPO	15.85	22.31	61.84
What is the coinsurance rate under the HDHP?	18.56	25.64	55.80
HDHP-Existing	33.85	21.24	44.91
HDHP-New	29.07	21.37	49.56
PPO	15.66	26.61	57.73
Do you get to keep HSA funds after the end of the year?	75.69	9.23	15.08
HDHP-Existing	96.73	1.38	1.90
HDHP-New	94.22	1.75	4.03
PPO	71.23	10.96	17.81
How much is $1000 worth in pre-tax dollars?	14.50	44.86	40.64
HDHP-Existing	16.93	31.78	51.30
HDHP-New	15.76	42.73	41.51
PPO	14.09	46.58	39.33

Table A1: Overall statistics here are weighted by plan density in the full sample.

	None	<1 hour	1-5 hours	6-10 hours	11-20 hours	>20 hours	Not sure
How much time do you expect to spend in the HDHP?	1.76%	5.99%	21.73%	17.40%	12.88%	24.92%	15.34%
HDHP-Existing	5.18	19.17	46.11	17.62	5.53	6.39	-
HDHP-New	3.50	14.71	40.81	22.24	11.21	7.53	-
PPO	1.17	3.52	16.83	16.83	13.89	28.96	18.79

	Understand, not concerned	Accept, but concerned	Don't like, no matter what
How do you feel about spending time managing your health plan?	14.82%	42.52%	42.65%
HDHP-Existing	39.03	32.64	28.32
HDHP-New	26.62	39.05	34.33
PPO	10.76	44.04	45.21

	Same	HDHP bigger	PPO bigger	Not sure
How do the provider networks of the two plans compare?	34.52%	6.04%	12.46%	46.98%
HDHP-Existing	41.28	6.74	2.76	49.22
HDHP-New	49.39	3.33	4.20	43.08
PPO	32.09	6.26	14.48	47.16

	Correct	Overestimate	Underestimate	Not sure
How much was spent on you and your dependents in 2011?	36.66%	29.81%	23.31%	10.22%
HDHP-Existing	41.97	35.75	16.41	5.87
HDHP-New	37.13	27.85	23.47	11.56
PPO	36.01	29.35	24.07	10.57

	Yes	No	Not sure
Do you think you will benefit/would have benefited from the HDHP in 2012?	16.49%	58.35%	25.16%
HDHP-Existing	56.65	23.83	19.52
HDHP-New	30.47	42.91	26.62
PPO	10.37	63.99	25.64

Table A2: Overall statistics here are weighted by plan density in the full sample.

B Appendix: Cost Model Setup and Estimation

This appendix describes the details of the cost model, which is summarized at a high-level in section 4.[36] 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.[37] 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 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.The level of precision we gain from the cost model leads to more credible estimates of the choice parameters of primary interest (e.g. risk preferences and information friction impacts).

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:

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

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:

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

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

[36] The model is similar to that used in Handel (2012).

[37] In the consumer choice model, this is mostly useful for estimating out-of-pocket expenditures in the $HDHP$ plan, since the PPO plan has essentially zero expenditures.

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 12, taking into account the trade off between cell size and precision. The minimum number of individuals in any cell is **XXXXXX** while almost all cells have over **XXXXX** members. Thus since there are four categories of claims, each individual can belong to one of approximately 10^4 or 10,000 combination of cells.

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

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

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

The above process generates a distribution of claims for each d and z but does not model correlations 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').[38] Here, G_{idt} is the distribution G_{dz} where $i \in I_{dz}$ at time t. Since correlations are modeled across d we pick the metric θ to group people into cells for the basis of determining correlations (we use the same cells that we use to determine group people for hospital and physician office visit claims). Denote these cells based on θ by z_θ. Then for each cell z_θ denote the empirical rank correlation between claims of type d and type d' by $\rho_{z_\theta}(d, d')$. Then, for a given individual i we determine the joint distribution of claims across D for year t, denoted $H_{it}(\cdot)$, by combining i's marginal distributions for all d at t using $\rho_{z_\theta}(d, d')$:

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

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

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

$$\Omega_j : C \to OOP_j$$

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

[38] 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 A-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 ensure 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 B-1 **TO BE ADDED** shows the results of one validation exercise for the $HDHP$ plan. 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 A-2 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

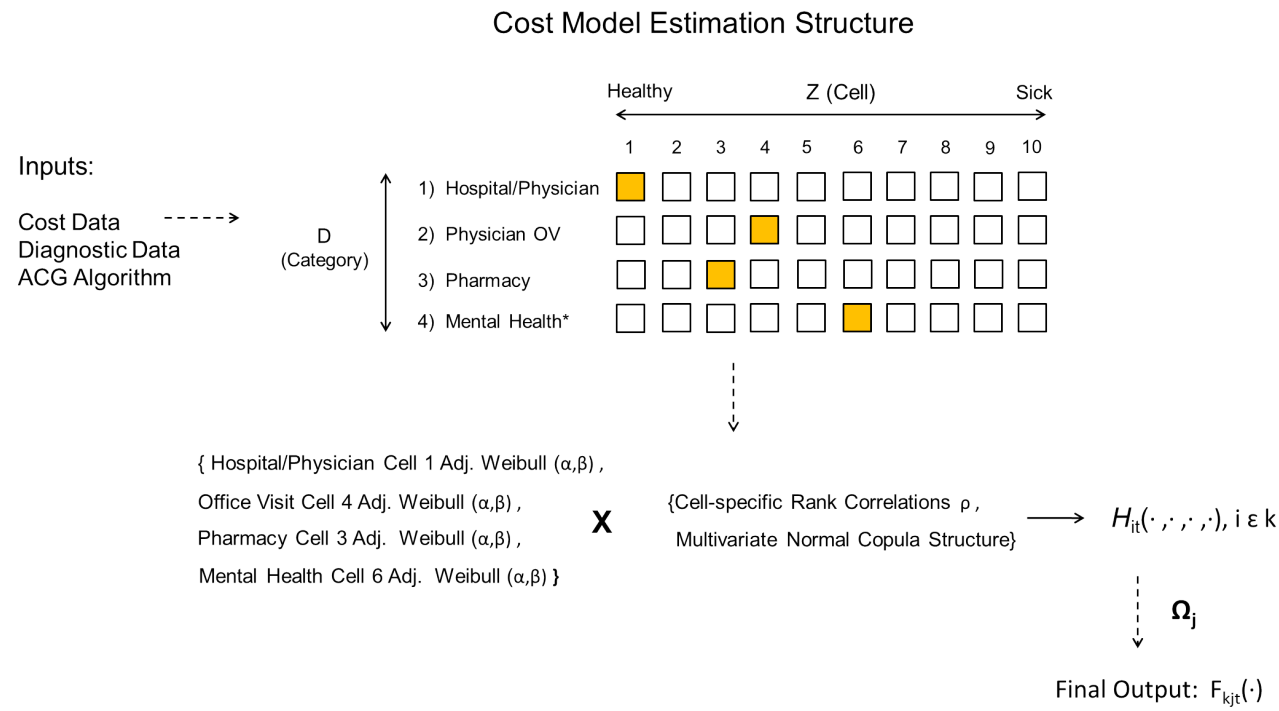


Figure B1: This figure outlines the primary steps of the cost model described in Online Appendix A. It moves from the initial inputs of cost data, diagnostic data, and the ACG algorithm to the final output F_{kjt} which is the family, plan, time specific distribution of out-of-pocket expenditures that 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}.

of claims for each cell z in claim category d, G_{dz}. To do this, we fit the empirical distribution of claims $G_{I_{dz}}$ to a Weibull distribution with a mass of values at 0. We use the Weibull distribution instead of the log-normal distribution, which is traditionally used to model medical expenditures, because we find that the log-normal distribution over-predicts large claims in the data while the Weibull does not. For each d and z the claims greater than zero are estimated with a maximum likelihood fit to the Weibull distribution:

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

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

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

Again, we use the notation $\hat{G_{iDt}}$ to represent the set of marginal distributions for i over the categories d: the distribution for each d depends on the cell z an individual i is in at t. We combine the distributions $\hat{G_{iDt}}$ for a given i and t into the joint distribution H_{it} using a Gaussian

copula method for the mapping Ψ. Intuitively, this amounts to assuming a parametric form for correlation across $\hat{G_{iDt}}$ equivalent to that from a standard normal distribution with correlations equal to empirical rank correlations $\rho_{z_{\theta_{it}}}(D, D')$ described in the previous section. Let $\Phi^{i}_{1|2|3|4}$ denote the standard multivariate normal distribution with pairwise correlations $\rho_{z_{\theta_{it}}}(D, D')$ for all pairings of the four claims categories D. Then an individual's joint distribution of non-zero claims is:

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

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

C Appendix: Choice Model Estimation

TO BE ADDED