HEALTH INSURANCE FOR "HUMANS":
INFORMATION FRICTIONS, PLAN CHOICE, AND CONSUMER WELFARE

Benjamin R. Handel
Jonathan T. Kolstad

ABSTRACT

Traditional models of insurance choice are predicated on fully informed and rational consumers protecting themselves from exposure to financial risk. In practice, choosing an insurance plan from a set of complex non-linear contracts is a complicated decision often made without full information on several potentially important dimensions. In this paper we combine new administrative data on health plan choices and claims with unique survey data on consumer information and other typically unobserved preference factors in order to separately identify risk preferences, information frictions, and perceived plan hassle costs. The administrative and survey data are linked at the individual level, allowing in-depth investigations of the links between these micro- foundations in both descriptive and choice-model based analyses. We find that consumers lack information on many important dimensions that they are typically assumed to understand, perceive high plan hassle costs, and make choices that depend on these frictions. Moreover, in the context of an expected utility model, including the additional frictions that we measure has direct implications for risk preference estimates, which are typically assumed to be the only source of persistent unobserved preference heterogeneity in such models. In our setting, we show that incorporating measures of these frictions leads to meaningful reductions in estimated consumer risk aversion. This result has both positive and normative implications since risk aversion generally has different welfare implications than information frictions. We assess the welfare impact of a counterfactual menu design and find that the welfare loss from risk exposure when additional frictions are not taken into account is more than double that when they are, illustrating the potential importance of our analysis for policy decisions.

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

Jonathan T. Kolstad
The Wharton School
University of Pennsylvania
306 Colonial Penn Center
3641 Locust Walk
Philadelphia, PA 19104
and NBER
jkolstad@wharton.upenn.edu

1 Introduction

In both employer-sponsored health insurance markets and the health insurance exchanges introduced as a part of national health reform consumers grapple with how to choose an insurance plan from a menu of options. As in the markets for other complex products, such as, e.g., cellular phone plans or financial investment vehicles, in health insurance markets real-world consumers may struggle to either obtain or process information in a way consistent with the *homo economicus* model typically used to study behavior in these settings. How consumers value different product attributes, what consumers know about those attributes, and how these preferences and information translate into choices is fundamental to market design and regulation, for health insurance and beyond. Without detailed knowledge of these micro-foundations it is difficult to precisely answer key policy questions such as which type of plans to allow insurers to offer and how those plans should be presented and priced.

Accordingly, there has been much recent empirical work that seeks to estimate micro-founded models of consumer insurance plan choice and then use those estimates for welfare analysis, in some cases for counterfactual market policies (see e.g. Bundorf et al. (2012), Cohen and Einav (2007), Carlin and Town (2009), Einav et al. (2010c), Einav et al. (2013), Abaluck and Gruber (2011), and Handel (2013)). One common aspect across these studies is their use of detailed administrative data on plan choices and risk realizations to identify crucial demand factors such as risk preferences and risk expectations.[1] These studies are typically unable to identify multiple unobserved preference factors apart from risk preferences because of the limitations of typical administrative data: the choices that consumers make, conditional on their risk expectations, are the primary instrument available. As a result, researchers use these observed choices to identify risk preferences, under assumptions that directly specify the roles of other unobserved choice factors, such as the information consumers have about available plan options.

While such assumptions are necessary given the data available, 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 key plan design aspects or even their own medical expenditure risk (see e.g. Kling et al. (2012),Ketcham et al. (2012) or Fang et al. (2008)). In addition, prior work such as Abaluck and Gruber (2011) and Barseghyan et al. (2012) has shown that consumers may exhibit decision-making biases even conditional on their information sets.[2] Finally, potentially important plan attributes such as time and hassle costs of actually using an insurance plan can differentiate even actuarially identical options but are typically unobserved.

If these foundations matter and are assumed away there are several key implications. First, in structural analyses where researchers are interested in quantifying specific choice foundations, (e.g. risk preferences) and using those estimates for counterfactual choice predictions, omitting relevant

---

[1]In health insurance, these studies occur in a variety of empirical contexts, ranging from large employer insurance markets to Medicare Part D prescription drug insurance.

[2]Grubb and Osborne (2013) find similar behavior in cellular phone markets, where consumers also chose from menus of potentially complex non-linear contracts.

unobserved factors will bias the conclusions drawn. Second, distinguishing between such choice factors can be important for welfare analysis, even in non-structural analyses such as Einav et al. (2010b) that model demand without specific assumptions on choice micro-foundations. In such frameworks, if unobserved preference factors are 'welfare-relevant' in the sense that they directly impact consumer welfare conditional on enrollment, then estimating demand is sufficient to conduct policy analyses; observed choices directly reflect relative 'ex post' plan valuations. If, however, unobserved factors such as consumer information or beliefs impact consumers choices, but not consumer welfare once enrolled, then neither reduced form demand curves nor structural analyses that omit such factors provide sufficient measures to conduct welfare analysis. 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. This distinction has been demonstrated theoretically (e.g. Spinnewijn (2012) and Bernheim and Rangel (2009)) though, to our knowledge, there is limited empirical work that makes the distinction between welfare-relevant and non-welfare-relevant choice factors.[3] This is due, at least in part, to the challenges to gathering data that identifies choice foundations beyond the standard model.

To overcome this empirical challenge, 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 other typically unobserved demand factors such as plan time and 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 from a menu of two plans, subsequent medical claims, demographics, and employment characteristics. The survey, administered electronically to a random sample of 4,500 employees soon after the open enrollment period, asks consumers simple questions designed to measure the information they possess on plan financial characteristics (e.g. deductible, co-insurance, OOP maximum), non-financial plan characteristics (e.g. provider network differences), and beliefs about their own total medical expenditure risk. In addition, we ask about the time and hassle costs of plan use that consumers have experienced and that consumers perceive for each plan option. The addition of rich individually-linked survey data to detailed administrative data adds multiple instruments that can be used to distinguish between risk preferences and other potentially important unobserved choice factors.[4]

We present several model-free descriptive analyses to illustrate the importance of information frictions and hassle costs for consumer choices. In our setting, consumers choose between two plan options: a PPO option with comprehensive risk protection and a high-deductible health plan

[3]Beshears et al. (2008) discuss potential ways to distinguish between revealed and normative preferences. In concurrent work, Baicker et al. (2012) studies medical care utilization with a welfare model that also implies a gap between the choices consumers make and the choices that maximize their welfare if fully informed.

[4]One potential downside to using survey data is that it relies on elicitations, rather than exogenous variation in administrative data, to identify these additional choice factors. While an "ideal" investigation of these factors would use only administrative data with exogenous variation on many dimensions, in practice this has not been done and seems unrealistic.

(HDHP) option with the same medical providers and treatments as the PPO, lower relative up front premiums, and larger relative risk exposure. First, before incorporating the linked survey data, we show that the choices consumers make suggest substantial risk aversion if risk aversion is the primary unobservable preference factor: in expectation, many consumers could have gained substantial monetary value by switching to the less comprehensive HDHP option without taking on substantial downside risk. Second, we investigate the correlations between answers to information-related survey questions and plan choices, conditional on realized costs, to illustrate that consumers that are relatively less informed about the HDHP option are less likely to choose that plan. A leading example comes from a simple question we ask concerning whether consumers know they can access the same medical providers and treatments in the two plans (they can). Approximately 20% of consumers incorrectly believe that the more financially comprehensive PPO plan grants greater medical access while 30% answer that they are "not sure" how the access compares. We show that these consumers are much more likely to choose the PPO relative to individuals who know that the plans grant exactly the same access. We illustrate similar implications for choices due to a lack of relative information on various plan financial characteristics and due to relative perceptions of hassle costs. Finally, we perform several analyses to alleviate concerns about the validity of the survey instrument including discussions of (i) reverse causality from experiential learning and (ii) confirmatory bias. Overall, our descriptive analyses suggest that information frictions and hassle cost perceptions matter for choices and that, if we omit these factors from our choice model, we will overestimate risk preferences in our setting due to the structure of plans and frictions present.

We study the importance of explicitly accounting for these additional frictions by estimating a series of structural choice models. These include (i) a baseline model, based just on administrative data, with risk preferences and health risk (ii) our primary model that adds information frictions and hassle costs measures derived from the linked survey and (iii) a types model that aggregates measures of information frictions into a one-dimensional information index. All models reflect expected utility maximizing, risk averse consumers while the models with additional frictions allow for consumers who may not have full information on relevant choice factors or may believe that one plan has higher hassle costs. Each model incorporates the output from a detailed ex ante cost model that predicts future health expenditure distributions at the time of plan choices. Comparison between the baseline model, which bears some similarity to those in the literature, and each model with additional frictions allows us to quantify the importance of choice frictions for consumer choices and how much risk preference estimates are biased by omitting these additional frictions from the analysis.[5]

Our estimates reveal the important role of the additional frictions. The baseline model, based on the administrative data alone, predicts substantial risk aversion, with a mean constant absolute risk aversion (CARA) coefficient of $1.60 \cdot 10^{-3}$. Framed in terms of a simple hypothetical gamble

---

[5]Since consumer inertia could be an important factor in our choice setting, the baseline model we emphasize also includes estimates of inertia identified in the administrative data by comparing the choices made by new employees to those made by existing employees. Our conclusions on the impact of including additional frictions for risk preference estimates are robust to the model of inertia used: as the estimates / model of inertia change, the implications of our information friction measures change but risk preference estimates in the full model are close to unchanged.

of similar scale, a consumer with this level of risk aversion would only be indifferent between not taking any action and taking on a gamble in which he gains $1000 with a 50% chance and loses $367 with a 50% chance. In other words, he would have to be paid a risk premium of roughly $633 in expectation to take on this risky bet. Incorporating measures of inertia, consumers are estimated to be less risk averse: the average one would be indifferent between no gamble and the same gamble that loses $812 with a 50% chance rather than $367.[6] Our primary model — incorporating information frictions — leads to lower estimates of risk aversion relative to both baseline models: in the full model with all frictions the consumer would be indifferent if the gamble included a 50% loss of $913, while in the types model this value is $924.

Focusing on specific frictions, we find the most influential frictions we measure are a lack of information about available medical providers / treatments and perceived time and hassle costs for the HDHP plan. For the former, a consumer who incorrectly believes that the PPO option grants greater medical access than the HDHP leaves, on average, $2,267 more on the table by choosing the PPO relative to a correctly informed consumer. The median consumer leaves $119 on the table per extra hour of perceived hassle cost in the HDHP relative to an otherwise similar consumer. Aggregating across all frictions we model, the average consumer leaves $1,694 on the table to choose to PPO over the HDHP relative to a fully informed consumer with zero perceived hassle costs. Without the linked survey data, these factors would be proxied for incorrectly by risk preference estimates, but once we include them the degree of estimated risk aversion is reduced.

We illustrate the welfare implications of these results by studying the impact of a counterfactual plan menu design that removes the PPO option from the choice set and forces all consumers to enroll in the high-deductible plan. We hold premiums and all other plan characteristics constant to focus on the key welfare distinction between risk preferences and information frictions. This permits a direct investigation of the welfare distinction without necessitating either a supply-side framework or a welfare model for information acquisition; since consumers are forced into the high-deductible plan, it is simply the intrinsic welfare of being enrolled in that plan that we care about.[7] A second advantage of our counterfactual is that the firm we study actually implemented this menu change after the period of our analysis. The exercise is thus relevant both to this large firm, as well as to other large employers and market designers (e.g. for health exchanges) considering similar menu design options.

Our analysis assumes that, even though information frictions impact choices, conditional on choosing a plan those frictions do not actually impact welfare. This implies, for example, that even if a consumer doesn't know that medical care access under both plans is identical, once forced to enroll in the HDHP this ex ante lack of information doesn't matter for welfare.[8] Nevertheless,

---

[6]This suggests that, in our setting, if one has just administrative data, incorporating inertia into the model matters a lot for risk preference estimates. In the recent literature mentioned earlier, people usually either model inertia explicitly (e.g. Handel (2013)) or study active choice settings (e.g. Einav et al. (2013).

[7]One extension might include a supply side with multiple plans, paired with information provision to reduce information frictions, permitting analysis of how these frictions impact market equilibrium and adverse selection.

[8]This same logic extends easily to other information frictions. On the other hand, time and hassle costs could have tangible welfare implications once enrolled. We examine a range of scenarios ranging from the (baseline) case where hassle costs are not welfare relevant (e.g. due to ex ante misperceptions or counterfactual improvements in

because omitting these choice frictions changes risk preference estimates, the welfare impact of different policies depends directly on accounting for these factors in the empirical choice model. Relative to the baseline case of risk neutrality, we find that the full model estimates, with lower risk aversion, imply an average welfare loss of \$62 per person from risk exposure, while the baseline model with (without) inertia implies a more than double \$148 (\$511) relative loss. We illustrate the implications of these results for a specific policy decision by viewing them in light of the fundamental tradeoff between risk protection and moral hazard inherent to optimal insurance design (see e.g. (Zeckhauser (1970)). Under the baseline model, with higher risk aversion, a price elasticity of demand for health care utilization of at least 0.280 would be necessary to justify the policy shift to the HDHP, while under the full model the elasticity would be 0.178.[9] Thus, for policymakers using an elasticity of .18 equivalent to that in the RAND Health Insurance Experiment (HIE) (Newhouse (1993)) the inclusion of information frictions in the model would be pivotal to their decision.

We note that all results presented here are specific to the large employer context that we study. From a theoretical perspective, incorporating information friction and hassle costs measures into typical insurance choice models could either increase or decrease the extent of estimated risk aversion. The direction of this effect will depend directly on the plans consumers can choose between and the relative information they have about each option. We illustrate here that the additional choice factors we study can matter for choice analysis, welfare analysis, and policy analysis, but the exact implications will depend on the specific context.

The paper proceeds as follows. Section 2 develops a conceptual theory of insurance choice. Section 3 describes the data, empirical setting and presents some descriptive analyses. Section 4 develops our empirical model of insurance choice. Section 5 presents results. Section 6 presents our welfare analysis of the counterfactual menu design we consider while 7 concludes.

2 Foundations of Choice in the Health Insurance Market

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 insurance plan decision depends 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 (or family), indexed by k, choosing health insurance plan j from a set of options $\mathcal{J}$. The consumer's utility for plan j is:

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

Here, W_k is wealth, P_{kj} is the premium facing individual k in plan j, and $f_{kj}(s|\psi_j, \mu_k)$ is the probability density of out-of-pocket expenditures in plan j for individual k. Out-of-pocket spending

---

plan design) to the case where they are fully relevant upon forced enrollment.

[9]These results assume zero marginal value of medical care foregone. If consumers value the care foregone at the high-deductible plan coinsurance rate, these elasticities are 0.407 and 0.258 respectively

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 ex ante health status.[10] Together, the terms of the plan and total spending distribution 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 it reduces the mean or variation of expected out-of-pocket spending and their willingness to pay for the latter is increasing in risk aversion. 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. Further, the model elements can be observed in widely available administrative data sets (e.g., expected expenditures for an individual and the plan options).

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 total price paid and the hospitals and doctors one can access. The introduction of Health Savings Accounts (HSA) and Flexible Spending Accounts (FSA) have introduced additional plan attributes not directly related to consumer risk protection. Plans can also have varying degrees of time and hassle costs linked to plan administration and logistics (e.g. dealing with medical bills). More generally, health insurance plans are differentiated products across a variety of dimensions beyond simple financial risk protection.

We extend the model to account for additional components of the choice problem that are not directly related to financial risk.[11] Plans differ by the network of physicians and hospitals available, the time and hassle costs associated with dealing with claims, and the tax benefits of linked financial accounts. Here, for exposition, we subsume these non-financial attributes with a plan-specific shifter $\pi_j(\psi_j, \mu_k, (1-t_k))$ that depends on plan design (ψ_j) and consumer type (μ_k) to reflect the fact that utility for these factors can depend on consumption of care and illness.[12] π_j also depends on an individual's marginal tax rate, to reflect the value of FSA and HSA contributions. Incorporating these features into the model utility from plan j for individual k yields:

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

---

[10]For the case of a family buying insurance, μ_k is a vector of health status types for all family members.

[11]The inclusion of these features in models of insurance choice is not new (see e.g. Ho (2009), Cutler et al. (2000)). However, measurement of these plan attributes, and preferences for them, has been difficult for researchers.

[12]In our empirical model, we model each of these non-financial attributes as a distinct factor. Here, π_j can be thought of as a utility model for each of these factors.

In this model, consumers still value plans as tools for risk protection. In addition, though, consumers can also be more willing to pay for a plan with different attributes, even if their distributions of financial losses from illness are identical in two plans.

Information Frictions in Plan Choice. In the description above, the choice of insurance plan relies entirely on individuals' risk preferences, their expenditure projections, and plan attributes. Importantly, the model above assumes that, when individuals make insurance choices, they can access and process the necessary information to make the correct decision under uncertainty. Accordingly, individual choices reflect real preferences for trading off premiums in exchange for shifts in either the distribution of out-of-pocket spending or non-financial attributes across different plans. This assumption is critical and underlies positive analysis of choice patterns throughout the literature on health insurance markets. Without this assumption, assessing welfare using revealed preference becomes more challenging (see e.g. Spinnewijn (2012) and Bernheim and Rangel (2009)).

There are 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. Without the assumption of full information, in the standard model where preferences are merely over financial risk the consumer might not know or understand the financial attributes that differentiate each plan, implying an inability to accurately forecast spending in each option. 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:

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

We assume that individuals observe each type of plan attribute with two types of error. The first is standard, mean zero, measurement error captured by ϵ. The second is an attribute specific shifter, δ, that captures information frictions in the model. Consumer choices no longer necessarily reflect the exact 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 plan utility is based on their beliefs about plan attributes and cost as follows:

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

From (3) we see how information frictions can impact the choice behavior of consumers in potentially important ways. Since both $\widehat{\psi_j}$ and $\widehat{\mu_k}$ enter the choice problem and impact the perceptions of (and

subsequent responses to) out-of-pocket expenditure risk, even if we observe the choices of individuals who optimize given their beliefs, we cannot necessarily recover key features of the model, such as risk preferences, with the standard model and typical administrative data. Similarly, if individuals are imperfectly informed about the non-financial attributes of the plan this will lead to choices that differ from what would have occurred with full information on the plan's network of physicians, true time and hassle costs, or a correct understanding of the tax benefits of plan features such as an HSA.

While choices may be affected by information frictions, these frictions may not impact true, welfare-relevant, utility conditional on enrolling in a given plan option (captured in equation (2)). For example, if a consumer believes that the providers available in-network in two plans differ, when they are in fact the same, this will impact choices but should not impact actual ex post consumer utility and welfare for one option relative to another. Thus, when information frictions impact choices, the standard model may (i) omit key choice foundations (ii) have biased estimates of the foundations estimated, such as risk preferences, and (iii) lead to biased assessments of the welfare impact of different market environments 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 measures of information and 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 information about plan attributes and key risk characteristics. The remainder of the paper focuses on developing an empirical model, related to equation (3), to assess the positive impact of information frictions on choice as well as the impact of including information frictions on welfare predictions for different counterfactual scenarios.

3 Data and Descriptive Analysis

We study health plan choice and utilization for the employees (and dependents) of a large self-insured employer with approximately 55,000 U.S. employees (in 2012) covering approximately 160,000 lives. We observe detailed administrative data with several primary components over the time period 2009-2012. First, we observe detailed characteristics of the health insurance choices that employees have in each year, as well as the choices that they ultimately make. Second, we observe the universe of line-by-line health care claims for all employees and their dependents in all plans. This includes detailed payment information, such as the total payment for a given service and the employee out-of-pocket payment, as well as diagnostic medical information that can be used to model health status. Finally, we observe demographic and linked choice information for each employee. For demographics, this includes, e.g., information on myriad job characteristics, income, age, and gender. For other choices, we observe, e.g., health savings account (HSA) participation and contributions, flexible spending account (FSA) participation and elections, and 401(k) contributions. These administrative data have similar components to that used recently in the literature

studying insurance provision at large self-insured firms (see e.g. Einav et al. (2013), Carlin and Town (2009), or Handel (2013)). The combination of these data with individually-linked survey data allow us to move beyond this recent work and study multiple additional micro-foundations and choice frictions that could impact both plan enrollment and consumer welfare.

The first column of Table 1 presents summary statistics for all employees present in all four years in the data from 2009-2012. There are 41,361 employees present in all four years, covering a total of 115,136 lives.[13] The employee population is heavily male (76.4%), young (49.7% $\leq$ 39 years)and high income (50.7% $\geq$ $125,000$) relative to the general population. 23% of employees are single, covering only themselves, with 19% covering a spouse only and 58% covering at least a spouse plus a dependent. Mean total medical expenditures for a family was $10,191$ in 2011. While the population we study is specific to our firm, implying the final numbers have limited external validity, 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.

Health Insurance Choices. Over the entire period 2009-2012, employees at the firm choose between two primary health insurance options a PPO option with generous first dollar coverage and a high-deductible health plan (HDHP) with a linked health savings account (HSA). We focus our analysis on the years 2011-2012 to match the time frame of our linked survey data.[14] The PPO option has had the largest share of employees over time, and had been the primary health insurance plan for many years prior to the introduction of the HDHP option in 2009. Since the HDHP introduction, the firm has promoted that financial benefits of that plan to employees in order to incentivize employees to economize on wasteful medical expenditures (while returning some of those savings in the process). For 2013, just past the end of our study period, the firm transitioned away from the PPO option and moved all employees previously enrolled there to the HDHP. Our counterfactual analysis in Section 6 studies the welfare implications of this menu change.

Table 2 compares the important characteristics of both plans. The PPO and HDHP have substantial differences in financial characteristics (e.g. premium, deductible, out-of-pocket maximum, HSA benefits) but, conditional on these financial elements, are identical on all other key features. Crucially, the HDHP offers access to the same set of in-network providers and the same medical treatments (at the same total cost), both key inputs into plan value. This allows us to model consumer welfare for enrolling in either of the plans as a function of financial characteristics only. On the financial dimension, the PPO option is the simpler and more comprehensive of the two options in terms of cost-sharing: it has no in-network deductible, no in-network coinsurance, and

---

[13]This sample is about 80% of the size of the mean number of employees present in *each* year from 2009-2012. We present descriptives for this 'full sample' as a baseline since this is the sample we use to estimate models with all employees, as described below. This sample also omits people who select the sparsely chosen HMO option that we exclude from the analysis.

[14]Depending on the location of the office within the U.S., a subset of employees could also choose a Health Maintenance Organization (HMO) option. Since approximately 5% of employees in the relevant locations choose this option (remaining steady over time) we exclude those who choose the HMO from our analysis and do not include the HMO option in our choice estimation.

Sample Demographics

	Full Sample	Survey Recip. (Weighted)	Survey Resp. (Weighted)
N -Employees	41,361	4500	1661
N_d - Emp.& Dep.	115,136	11,690	4,584
2011 PPO%	88.8	89.6	88.7
2012 PPO%	82.7	83.0	81.6
2011 HDHP %	11.2	10.4	11.3
2012 HDHP %	17.3	17.0	18.4
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%
$\geq$60	1.3%	0.9%	1.2%
Income			
Tier 1 ($<$ \$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 1: 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 re-weighted statistics for the second and third columns, as described in the text.

no in-network out-of-pocket maximum.[15] Alternatively, the HDHP has a substantial deductible equal to $1500 for individuals, $3000 for a couple (or parent and one child), and $3750 for a family. In that plan, once an employee spends an amount in excess of the deductible, he must then pay co-insurance of 10% of allowed costs for in-network providers and 30% for out-of-network providers until his total spending exceeds the out-of-pocket maximum — $2500 for individuals, $5000 for a couple, and $6250 for a family — at which point all expenditures are paid by the insurer.

The PPO plan charges no up front premium while the HDHP provides an up-front subsidy equal to $1500 for an individual, $3000 for a couple and $3750 for a family. This subsidy should be interpreted as the primary premium for the PPO relative to the HDHP. The HDHP subsidy is deposited into the health savings account (HSA) linked to that plan and, thus, can be used for medical expenditures on a pre-tax basis in both the short-run and the long-run. If employees want to use the subsidy for non-medical expenditures at any point in their lives, they can do so on a post-tax basis.[16] The linked HSA has the potential to provide additional value to the employee, above and beyond the subsidy. Employees can contribute their own-funds pre-tax to the HSA, up to a maximum of $3150 for individuals and $6250 for all others (gross of the subsidy). As with the subsidy, these incremental contributions can be used pre-tax for medical expenditures at any point in one's lifetime. Finally, in addition to the pre-tax benefits for medical expenditures, all HSA funds can be invested in a pre-tax manner over time, providing similar benefits to those of a 401(k) (or related investment vehicle).

Figure 1 depicts the financial returns to selecting the HDHP option relative to the PPO option for an employee in the family tier.[17] The x-axis plots realized total health expenditures (insurer + insuree) and the y-axis plots the financial returns for the HDHP relative to the PPO as a function of those total expenditures. At low levels of health care cost, the HDHP provides a relatively large financial return: when total expenditures are 0 this equals to the value of the subsidy, $3,750, plus any value from incremental HSA contributions. As total spending increases, these gains are reduced dollar-for-dollar before the deductible is met. Once total spending surpasses the deductible the family pays the coinsurance rate of 10% for incremental expenditures, diminishing the slope of the loss in the HDHP relative to the PPO as spending increases. Once out-of-pocket expending reaches the out-of-pocket maximum, relative value of the HDHP does not change for incremental total expenditures. The figure demonstrates that there is a unique level of expenditure above which the PPO plan is valuable *ex post* relative to the HDHP. Furthermore, the maximum financial loss from choosing the HDHP is $2,500 (assuming no valuable incremental HSA contributions).[18] Thus, for a family, the range of potential ex-post relative value for the HDHP spans $[-\$2,500, +\$3,750]$.[19]

---

[15]In the PPO employees have very limited spending for out-of-network expenditures as long as total charges don't exceed those from comparable in-network providers (exact characteristics are given in the table). Further, only approximately 4% of total expenditures are out-of-network.

[16]If they use these funds before 65 for non-medical expenditures, they pay an additional tax penalty of 10%.

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

[18]In a series of focus groups we conducted at the firm, the true magnitude of the maximum loss was particularly surprising to employees: many thought that the maximum financial loss in the HDHP would be larger. This underscores the complexity required to actually determine the money one could lose upon becoming sick.

[19]This range shifts upward by a constant amount if consumers derive value from incremental HSA contributions. The relative value range for an individual / couple equals the family bounds multiplied by 0.4 (0.8).

Health Plan Characteristics
Family Tier

	PPO	HDHP
Premium	0	0
Health Savings Account (HSA)	No	Yes
HSA Subsidy	-	$3,750*
Max. HSA Contribution	-	$6,250**
Deductible	0	$3,750*
Coinsurance (IN)	0%	10%
Coinsurance (OUT)	20%	30%
Out-of-Pocket Max.	0**	$6,250*
Provider Network	Same as HDHP	Same as PPO

* Single employees (couples) have value equal to .4 (.8) of family tier
**Single employees have a maximum of of$3,100 is max. contribution
while those over 55 can contribute an extra $1,000
***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 2: 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. The numbers in the main table are presented for the family tier (the majority of employees) though we also note the levels for single employees and couples below the main table.

Extending the analysis in Figure 1, we can compute the share of employees whose total medical expenditures were below the break-even point in 2011, determining those who would have been *ex post* better off in the HDHP. If we assume that all employees contribute the maximum amount to their HSA, and thus realize the maximum possible tax benefits, then 73% of employees would have been better off ex post in the HDHP in 2011. If we assume consumers make 50% of the maximum possible incremental HSA contributions (close to what is observed in the data) then 60% of employees would have been better off in the HDHP. If employees don't add any incremental HSA funds, then 35% would be ex post better off in the HDHP.[20]

Despite the potential value that the HDHP provides for consumers, relatively few choose that plan. As Table 1 reveals, in 2011 11.2% of employees in the full sample choose the HDHP, while in 2012 17.3% did. The actual choice percentages are much lower than the ex post optimal percentages just described. While some of this difference could be due to the difference between expected value ex ante and realized value ex post, since the downside loss in the HDHP is limited (and actually

---

[20]This analysis assumes a 35% marginal tax rate on income, near the average in the population.
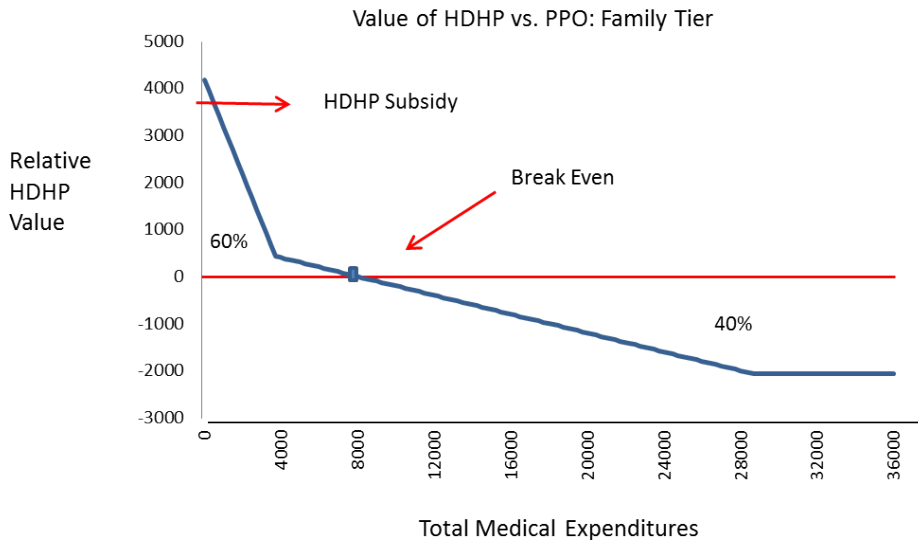
Figure 1: Description of HDHP financial value relative to the PPO in 2012, for the family tier, as a function of total medical expenditures. This chart assumes that employees contribute 50% of the maximum possible incremental amount to their HSA, near the median in the population. 60% of *all* employees break even, given their respective coverage tiers.

smaller than the potential upside) this is unlikely to explain much of this gap.[21]

This simple comparison suggests that consumers are choosing the PPO plan more than they 'should' from either an ex post perspective or from a risk-neutral ex ante perspective. An obvious reason for this could that consumers are risk-averse and value risk protection. Accordingly, the standard approach in the structural empirical literature would rationalize the observed choices by allowing for risk-averse consumers, with respect to financial risk.

Given the actual choices, estimating a model where risk-aversion and health risk are the primary choice drivers yields very high estimates of risk aversion (see results in Section 5). This should not be surprising, given that consumers have limited financial downside risk in the HDHP while there are equally large potential gains. Especially in light of the fact that about half of the employees in our sample earn over $125,000$, high risk aversion with respect to purely financial risk seems to be an unsatisfactory explanation for the low proportion of employees choosing the HDHP.[22] This low proportion could, however, also result from other factors that should matter for consumer choice in insurance markets, such as (i) a lack of information on plan features (ii) a lack of information on the distribution of possible total medical expenditures (iii) beliefs about non-financial attributes of the plan (i.e., time/hassle costs, physician networks, etc.) or (iv) actual differences in non-financial attributes of the plans (e.g., time/hassle costs). The remainder of the paper focuses on understanding which of these potential alternative micro-foundations can help explain observed choice behavior, as well as the differential welfare implications for those foundations relative to the

---

[21]For example, the percentage better off in the HDHP taking ex ante out-of-pocket expected values, and assuming risk-neutrality, is not very different than the ex post percentages.

[22]In our empirical model section, Section 4, we discuss in more detail the important distinction between classical risk aversion with respect to financial risk, and risk aversion with respect to informational issues.

traditional explanation of risk aversion.

Survey Data and Design. In order to measure information frictions and beliefs about non-financial plan attributes (such as time and hassle costs), we developed a survey instrument. In this section we discuss the key features of the survey as it pertains to our main analysis. Appendix A contains a more detailed discussion of the survey questions and methodology.

Our survey instrument was designed in conjunction with the both the Human Resources department and the Marketing and Communications department at the employer we study. The survey was administered by the firm's insurance administrator, a large private insurer, using clear and simple to navigate online format (see Appendix A for screen shots). The insurance administrator released the survey early in the calender year of 2012, and it remained opened for a period of two weeks, with reminders sent to the recipients just before the end of that period. The survey contained approximately thirty multiple choices questions. No incentive was given in the form of money or a prize to induce response. The survey was sent to 4,500 employees total, coming from three equal sized groups defined as (i) employees enrolled in the HDHP plan for both 2011 and 2012 ('incumbents') (ii) new HDHP enrollees in 2012 (almost exclusively people who switched from the PPO), and (iii) those in the PPO plan in both 2011 and 2012.[23] 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, implying an average overall response rate of 38%.

The three survey cohorts were specifically designed to over-sample the HDHP population relative to the PPO population in order to assure enough sample size for the former and ensure sufficient statistical power. In our primary analysis, we re-weight both the survey recipients and survey respondents to reflect that actual plan choice composition in the market. This re-weighting procedure follows the econometric literature on re-weighting, which advocates re-weighting based on the dimension of explicit oversampling (in our case plan choice). For a further discussion, see e.g. Solon et al. (2013) or Manski and Lerman (1977). Throughout our analysis, when we refer to our 'primary sample', we mean this re-weighted sample of survey respondents (or recipients when relevant).[24]

The last two columns of Table 1 present summary statistics for the randomly selected survey recipients as well as the well as the total survey respondents (both re-weighted as described above) and compares those samples to the full sample described in the first column. The different populations are, on the whole, quite similar, mitigating sample selection concerns for the survey respondents sample. Comparing the survey respondents to both the recipients and to the full population reveals that the populations are very similar in terms of age, gender, income and family size. The average spending is slightly higher among the respondents compared to the overall population,

---

[23]Very few employees enroll in the HDHP in 2011 and switch to the PPO in 2012.

[24]To implement the weighting in the most transparent manner, we use 'block' re-sampling where we construct a pseudo-population that re-uses the entire under-sampled group (the PPO cohort) K times where K is an integer that achieves the minimum distance between the choice proportions in the pseudo-population and the choice proportions in the actual full population. This methodology is simple to implement, and can be easily integrated with the block bootstrapping methodology used for estimating parameter standard errors in our econometric specifications.

but, comparing spending at different points in the distribution, this appears to be a small effect that is driven by higher spending in the upper tail of the cost distribution for respondents, rather than systematically higher spending across this distribution.[25][26]

We designed the survey to contain only multiple choice questions in order to have a simple format where we could clearly interpret question answers.[27] Each multiple choice question was motivated by our desire to learn about a specific dimension of consumer information, experience, or decision-making as described in our model in Section 2. Tables 3 and 4 summarize the primary questions used in our analysis and the responses from the survey population, broken down by cohort. The questions focus on four major areas of the benefits choice. The first targeted area assesses knowledge of the financial features of benefit design in the HDHP. These questions target information frictions directly as they ask respondents to correctly answer questions about key features of the HDHP. Each respondent was asked to correctly identify the deductible, coinsurance rate, out-of-pocket maximum, HSA subsidy level and tax benefits for HSA contributions from a set of options.[28] The second set of questions focused on a related source 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). The survey also asked a set of questions to determine whether respondents 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. These included questions about the time and resources required to manage both the HSA and the 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 asked questions about preferences for hassle in the HDHP. Finally, we asked a set of questions to ascertain the amount of effort that went into an employee's choice, the clarity of their beliefs about the plans, and their satisfaction with their choice.

Frictions: Descriptive Evidence. Before turning to our formal choice model, we present some descriptive results from our combined survey and administrative data to demonstrate the potential importance of the frictions we study. There are clear patterns in the raw survey responses that are consistent with limited information, as well as time and hassle costs. Furthermore, answers to

---

[25]Of course, the respondents could differ on unobservable dimensions (such as knowledge or degree of interaction with health benefits). However, if consumers who choose to answer are more well-informed than average, our results should reflect lower bounds on the impact of information frictions.

[26]We note that the survey recipients were selected at random from the entire population *after* removing a few thousand executive and top-level employees from the potential recipient pool. As a result, the recipient pool is slightly younger, slightly lower income, a little more likely to be single, and have slightly lower health care spending.

[27]We considered, e.g., including some belief elicitation or risk preference elicitation questions, but ultimately, together with the firm's Human Resources group, concluded we could best achieve our goals through transparent, information-based questions.

[28]Throughout the survey, much of our focus is on consumer information about and experience with the HDHP. An implicit assumption is that consumers have similar information about the simpler PPO option, and that, consequently, their answers to survey questions about the HDHP represent the relative difference in information about the HDHP and PPO. This could be thought of as assuming the everyone has close to 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. This assumption is supported by the questions we do ask consumers about the PPO.

Question	Correct	Incorrect	Not Sure
(1) 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
(2) 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
(3) 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
(4) 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
(5) 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
(6) 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 3: Responses to Plan Financial Characteristics Survey Questions. Exact wording of questions and answers in Appendix A.

some survey questions have a strong gradient with respect to actual plan choices made, even after conditioning on measures of health risk.

Table 3 describes consumer responses to questions that target knowledge of health plan financial characteristics. A (slim) majority of employees who were enrolled in the HDHP were able to correctly identify their deductible in that plan. Only slightly more than 20% of employees who enrolled in the PPO could identify the deductible for the HDHP choice option. In fact, more PPO enrollees answered incorrectly than correctly, though the majority were "not sure". A similar pattern holds for the questions asking about the post-deductible coinsurance rate and the out-of-pocket maximum in the HDHP, though fewer respondents have information on these characteristics, relative to the deductible. Approximately 70% of HDHP enrollees know the premium difference between the two plans, linked to the HSA subsidy, while only 20% of PPO enrollees do. Almost all HDHP enrollees know that HSA funds can be rolled over from year to year, while approximately

Question

	Same	HDHP bigger	PPO bigger	Not sure
(7) 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

	None	<1 hour	1-5 hours	6-10 hours	11-20 hours	>20 hours	Not sure
(8) 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
... in the PPO?							
PPO	15.85	29.75	29.16	11.35	2.94	4.11	6.85

	Understand, not concerned	Accept, but concerned	Don't like, no matter what
(9) 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

	Correct	Overestimate	Underestimate	Not sure
(10) 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

	Very confident	Somewhat confident	Not confident
(11) How confident are you in this estimate?	35.85%	43.90%	20.25%
HDHP-Existing	38.34	49.22	12.44
HDHP-New	30.11	46.13	23.77
PPO	36.20	43.05	20.74

	Yes	No	Not sure
(12) 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 4: Responses to Plan Non-Financial Characteristics, Hassle Cost and Medical Expenditure Survey Questions. Exact wording of questions and answers in Appendix A.

75% of PPO enrollees do. The answers to this question suggest that there is real information content in the survey question answers, as most PPO enrollees can answer this simple question regarding the HDHP correctly (rather than, e.g. "not sure"). Another pattern from Table 3 is that existing HDHP enrollees (enrolled in that plan for at least one year prior to 2012) have very similar answer proportions to new HDHP enrollees, who just signed up for that plan just before the survey. This suggests that experiential learning is not substantial, though without a formal model this conclusion should be viewed with some caution.

Table 4 presents the respondent answers to questions about non-financial plan characteristics. The first question asks respondents how the doctors and medical services that can be accessed in-network compare across the two plans. Recall that the networks are in fact, identical on all dimensions for the two plans. If consumers believe that one plan provides access to higher quality doctors, or a greater range of medical services, this could have a significant impact on their plan choices, even though we know that this should not impact their relative welfare between the two plan options, conditional on actually enrolling in either plan. 49% of incumbent HDHP enrollees, 41% of new HDHP enrollees, and 32% of PPO enrollees understand that one can access the same physicians in network in both plans. For both HDHP enrollee groups, almost all other answers are "not sure". 15% of PPO enrollees (who comprise most of the overall population) believe that the PPO provides greater access to physicians, 6% believe the reverse, and the remaining consumers are "not sure". 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 underscores the role of information frictions.

To better understand how important information about provider access is for explaining choices, Figure 2 studies plan choices as a function of the question answers. The left panel presents the share of enrollees in the HDHP conditional on their answers to this question. It is clear that those who understood that medical access was the same were far more likely to select the HDHP: 23% chose that plan, compared to 6% among those reporting the PPO had a larger network and 17% among those answering "not sure"). The right panel gives a sense of whether this relationship is caused by an underlying correlation between question answers and medical expenditures: it presents the optimal ex post choice based on actual 2011 expenditures. The figure indicates that a similar proportion of consumers should choose the HDHP across the survey question answer groups (between 30-40% with no incremental HSA contributions). This implies that the gap between the proportion of people who *should* choose the HDHP and those who *actually* do is much larger for those consumers who believe the PPO provides access to more physicians.

The second and third questions in Table 4 ask about consumers' expectations of and preferences for time and hassle costs stemming from plan administration and logistics (e.g., dealing with medical bills).[29] The hassle of dealing with paying for medical expenditures directly and being reimbursed is a potentially important non-financial attribute of the HDHP that might impact choice. The question on time and hassle cost expectations had 7 multiple choice options, ranging from "none"

---

[29]The actual question asked to employees is presented in Appendix A and is carefully worded so as to define what we refer to as time and hassle costs.
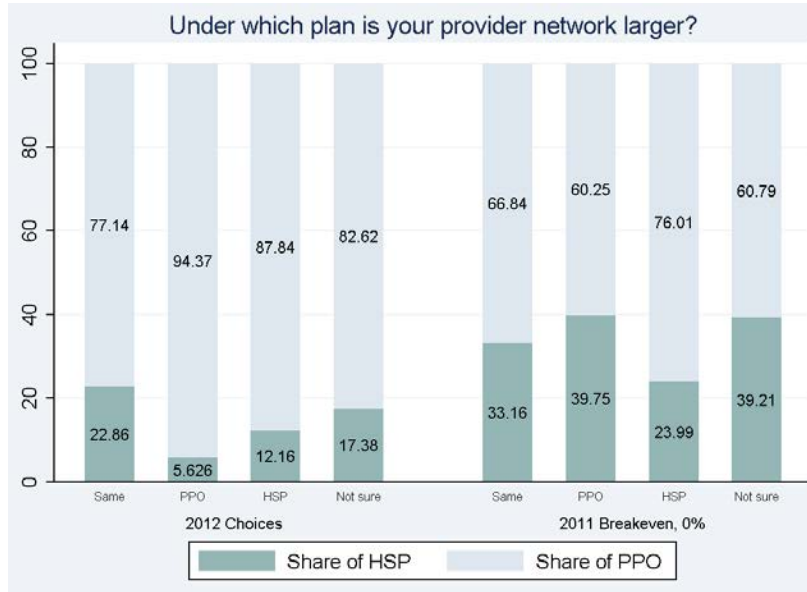
Figure 2: Actual versus Predicted Plan Choices by Knowledge of Plan Provider Networks

to "> 20 hours" ("not sure" was also an option). The results point to a substantial difference in perception of the time required to deal with the HDHP among those enrolled in the HDHP compared to PPO enrollees. For example, 29% of PPO enrollees answer that they would expect to spend more than 20 hours on HDHP plan administration and logistics, while only 6% and 8% do in the two HDHP cohorts. This is despite the fact that only 4% of PPO enrollees believe that that plan leads to "> 20 hours" in time / hassle costs. It is interesting to note that new HDHP enrollees have quite similar beliefs about time and hassle costs as incumbent enrollees who had already experienced the plan suggesting that the difference between HDHP and PPO enrollees is not due only to experience with the HDHP plan. The third panel in Table 4 demonstrates a strong relationship between plan choice and how accepting consumers are of the time required to deal with the plan hassle costs. For example, only 11% of PPO enrollees report not being concerned that they may need to spend time managing health care costs compared to 39% of existing HDHP enrollees.

Figure 3 studies plan choices as a function of time and hassle cost perceptions. There is a strong relationship between expected time/hassle costs and plan choices: as projected costs increase, consumers are much less likely to choose the HDHP. For example, 37.2% of consumers who expected to spend 1-5 hours on plan administration and logistics in the HDHP choose that plan, while only 5.1% of those who expect to spend > 20 hours on these activities choose that plan. We note that our measures of expected time and hassle costs could represent multiple micro-foundations. For example, an individual could expect to have higher time and hassle costs because they have higher medical utilization in general, or, their expected time spent, conditional on medical utilization, could be higher because of heterogeneity in ability to manage and navigate complex financial products.[30]

---

[30]Additionally, individuals could lack information on potential HDHP time and hassle costs. In this event, this friction represents both a tangible friction that leads to real welfare losses, and an information friction that may not

The right panel of Figure 3 reveals that that the relationship between plan choices and projected time/hassle costs is due in part, but not fully, to correlation between expected time and hassle costs and medical utilization. It presents the optimal ex post based on actual 2011 expenditures. The figure indicates that those who expect to have lower hassle costs also have lower expenditures and, thus, ignoring utility from those time/hassle costs, should choose the HDHP in higher proportions. However, the gap between these ex post optimal choices and actual choices becomes larger as expected time and hassle costs do, suggesting that differences in perceived time/hassle costs are only due in part to differences in medical utilization.
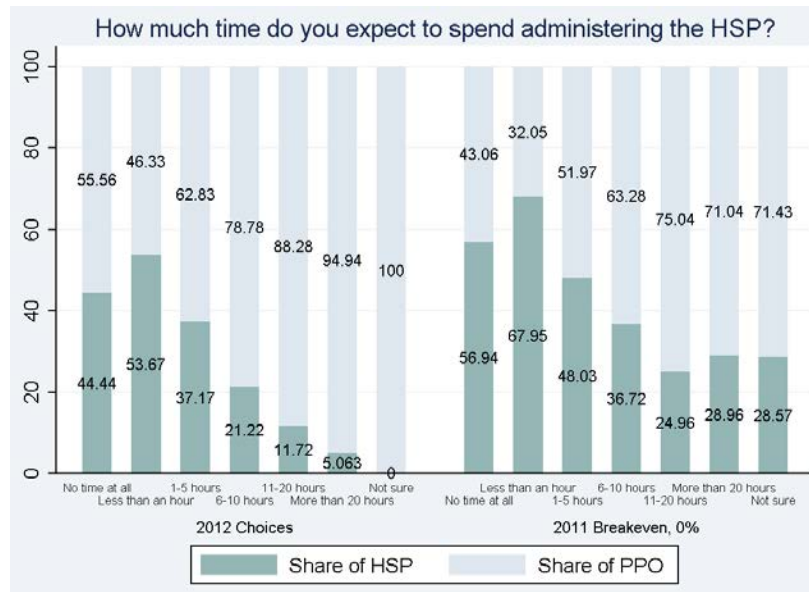

Figure 3: Actual versus predicted choices as a function of time and hassle cost perceptions.

Table 4 also presents the responses to questions asking about knowledge of total medical expenditures and knowledge of the tax benefits provided by a health savings account (HSA). In order to understand out-of-pocket expenditure risk in the HDHP, it is necessary to understand total potential medical charges as well as plan characteristics such as deductible and coinsurance. We ask consumers to identify their amount of total medical spending for the calender year 2011 (which had just ended at the time of the survey) and compare their answers to their actual total spending in that year. Consumers chose between the multiple choice options of $0-500, $501-2,500, $2,501-5,000, $5,000-10,000 and more than $10,000. The table presents the statistics for whether consumers overestimate, underestimate, or correctly guess their expenditures for the past year. Overall, the proportions in each of these buckets does not change much by cohort. Across the three cohorts, 36-42% answer the question correctly, 29-36% overestimate their past expenditures, and 17-24% underestimate them. 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 respondents in each cohort reply they are somewhat or very confident in their estimate. Thus, while it appears

_____

be welfare relevant. We return to this distinction in detail in our modeling and welfare analysis.

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 recognize this lack of understanding.

It is also important to understand correlation patterns in the answers to these questions. If survey responses are highly correlated across a given subset of questions, this could suggest that there are certain 'types' of consumers who have similar information content and choice frictions across these questions. Tables D1 and D2 in Appendix D present the full correlation matrix for the responses to our primary questions of interest. Table D1 studies correlations between the responses to the questions on plan financial characteristics presented in Table 3. The correlation between these answers are fairly high: for example, the correlation between knowing the deductible and knowing the post-deductible coinsurance rate is 0.35, while the correlation between knowing the deductible and knowing the HDHP subsidy is 0.45.[31] The correlation between answering "not sure" for deductible and "not sure" about the subsidy is 0.435. As a result of these patterns, in our upcoming empirical analysis we create one measure for information on plan financial characteristics that subsumes these highly correlated responses.

The degree of correlation is lower between responses to other questions, as shown in Table D2. For example, the correlation between expected time spent on plan administration and correctly knowing that the provider networks are the same for both plans is only 0.027. Additionally, the correlations between knowledge of plan financial characteristics and all other friction measures are relatively low, generally falling between -0.1 and 0.1. This suggests that there is meaningful multi-dimensional heterogeneity across these frictions, and that modeling them in a disaggregated manner could be fruitful. In our upcoming empirical analysis, we examine several specifications, ranging from a disaggregated specification that includes most friction measures as distinct variables to a types specification that develops a one-dimensional information index for consumers.

Survey Data: Discussion. As noted earlier, we believe that detailed survey data, linked to rich administrative data, can provide meaningful insights about information frictions and hassle costs, especially given that these factors are quite difficult to measure with administrative data alone. Here, we address two potential concerns related to the use of survey data. First, we discuss the possibility of *confirmation bias* whereby consumers who enroll in a certain health plan are more likely to choose the answers that favor the attractiveness of that plan, validating their recent choices (see e.g. Rabin (1998) for a richer discussion). Second, we discuss the possibility of experiential learning for consumers and the role that that plays in our analysis.

Our survey was administered at the beginning of the calender year 2012, a few months after the open enrollment period in November 2011. Given that consumers had already made their plan choices at the time of the survey, confirmation bias in survey responses would lead to consumers who select the HDHP (PPO) choosing answers that "confirm" or validate their choices. For example, someone who chose the PPO might answer that they believe that plan has access to more physicians in-network due to confirmation bias. We note that confirmation bias *does not* have anything to do

---

[31]Here, the variables are represented as binary variables, so a given correlation measure represents the correlation between two binary variables.

with search for information: if consumers who chose the HDHP were more likely to do research on the HDHP, and health plans in general, there is no issue since information set at the time of plan choice is exactly what we aim to capture.

While we can not rule out confirmation bias in any formal sense, there are several pieces of evidence that suggest it is not a particularly strong factor for our analysis. First, for many information-related questions, such as those on plan financial characteristics, PPO enrollees are more likely to answer "not sure" relative to HDHP enrollees: both groups are similarly likely to answer these questions incorrectly. "Not sure" suggests a lack of knowledge, but does not suggest validation of the PPO choice. Furthermore, for many more factually based questions (e.g., what was the deductible) it is not obvious that one answer is more preferential to a specific plan. Second, there is meaningful variation across questions in the proportion of consumers choosing answers that are favorable to the plan they chose. For example, 71% of PPO enrollees know that you can roll over HSA funds, an answer that is favorable to the HDHP plan, while only 6% believe the HDHP plan provides access to more doctors. Of course, confirmation bias could imply a shift in responses relative to some baseline information level that could differ across questions, but this evidence suggests that consumers are not blindly answering questions in order to validate choices.

Finally, multi-dimensional heterogeneity in frictions for consumers is suggestive of nuanced and informative answers. Table D2 in Appendix D reveals limited pairwise correlations between an aggregated measure of plan financial characteristic knowledge, knowledge about provider networks, expected time and hassle costs, and knowledge of own past medical expenditures. This suggests that if confirmation bias were present, it would have to manifest on different dimensions for different consumers, which we believe is less likely than the case where it is present on similar dimensions across consumers.

Experiential learning, where consumers learn about a plan while being enrolled, is not an issue for us if this learning occurred prior to or during the open enrollment period in 2011. We use the survey measures as proxies for information and expectations at the time of plan choice. Experiential learning could be an issue if consumers learn about a plan after they choose it during open enrollment but before the survey is administered near the beginning of 2012. This would lead to these consumers having more information than they had at the time of plan choice. It is also possible that in the period between open enrollment in November and the beginning of 2012 that consumers forget information that they knew when choosing between plans. While we cannot rule this out as a potential source of bias in our analysis, there are several reassuring features of our data and environment that make it unlikely this is a major factor. First, and most importantly, the survey was conducted near the beginning of 2012, indicating that experiential learning would have had to occur mostly before new HDHP enrollees had marked experience with that plan. Further, consumers who had already been enrolled in the same plan the year before (either the PPO or HDHP) are unlikely to have marked incremental learning in this short time period. Second, it is unlikely that consumers would forget simple pieces of information (such as that there are identical provider networks) over a short time period. We cannot rule out that consumers forget information about more complex objects (like a plan out-of-pocket maximum), though our multiple choice

answer format facilitates recall. Additionally, the strong positive correlations for answers to plan financial characteristic questions (shown in Table D1) suggests that information on these more complex dimensions is more of an 'all or nothing' proposition. Thus if consumes forgot information on these features over the short period between open enrollment and the survey administration, they would have had to forgotten most or all information they knew. Finally, we don't believe that the questions about time and hassle cost expectations should be markedly impacted by experiential learning during this short interim period.

4   Empirical Framework

The analysis in the previous section provides evidence that information frictions are present for a variety of key choice dimensions and are correlated with consumers' health plan choices in a manner that implies more informed consumers are choosing plans that provide them more value. In this section, we develop a series of models that quantify the impact of information frictions, perceived hassle costs, and risk preferences on health plan choices.

In order to illustrate how the inclusion of information friction and hassle cost measures impact risk preference estimates, we start with a 'baseline' model that includes just health risk, risk preferences, and health plan characteristics. We then add measures of information frictions and hassle costs derived from the linked survey in three ways. We first estimate incremental models that add one friction to the baseline model. We then incorporate all frictions in a 'full' model. Finally, we estimate a 'types' model that includes all information frictions aggregated into a one-dimensional index. In addition to the specific estimates of risk preferences and choice frictions in the 'full' and 'types' models, which may be of intrinsic interest, the structural approach allows us to study how risk preference estimates are impacted by including additional factors that link to plan choices. The distinction between choices based on risk preferences and choices based on information frictions and perceived hassle costs is crucial to the welfare analysis discussed in section 6. While risk preferences impact both choices and welfare, a lack of information may be relevant for choices *given a menu of options* but may not impact actual welfare *conditional on enrollment in an option*.[32]

Baseline Choice Model. The baseline model studies expected utility maximizing families who make active (non-inertial) choices and are fully informed about all health plan options. Consumer choices depend on (i) *ex ante* cost risk (ii) risk preferences and (iii) an idiosyncratic mean zero preference shock. We describe the baseline choice framework here *conditional* on our ex ante cost projections, which are estimated in a separate detailed medical cost model described later in this section and do not vary with the choice model specification. The model presented is the empirical analog to equation (1) in Section 2.

---

[32]We note that, as discussed in Einav et al. (2010b), for certain policy questions (e.g. pricing policies) a structural approach is not necessary to conduct welfare analysis. The same is true for welfare analysis that distinguishes between 'welfare-relevant' and 'non-welfare-relevant' choice factors, given appropriate available data, such as those discussed here. There are many questions, such as those involving counterfactual insurance plan designs, that do require a model of micro-foundations such as that developed here.

Denote the family-plan specific distributions of out-of-pocket health expenditures output by the cost model as $F_{kj}(\cdot)$. Here, $k \in K$ is a family unit, $j \in J$ is one of the two health plan options available at the firm in 2012. The baseline model assumes that families' beliefs about their out-of-pocket expenditures conform to $F_{kj}(\cdot)$. Each family has latent utility U_{kj} for each plan and chooses the plan j that maximizes U_{kj}. We assume that U_{kj} has the following von Neumann-Morgenstern (vNM) expected utility formulation:

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

Here, $u_k(\cdot)$ is the vNM utility index and s is a realization of out-of-pocket medical expenses from $F_{kj}(\cdot)$. W_k denotes family-specific wealth and x_{kj} represents consumption in a given state of the world (defined below). P_{kj} is the family-time specific premium for plan j. Formally, in our setting we define the premium $P_{k,HDHP}$ as:

$$P_{k,HDHP} = -(HSA_k^S + \tau_k HSA_k^C)$$

HSA_k^S is the firm's subsidy to each employee's health savings account (HSA) when they enroll in the HDHP. This is deterministic conditional on the number of dependents being covered (discussed in Section 3). HSA_k^C is the incremental contribution a family makes to the HSA, on top of HSA_k^S, when they sign up for the HDHP. The value of these contributions is equivalent to the value of pre-tax dollars relative to post-tax dollars, and thus depends on marginal tax rate τ_k.[33] Empirically, we model HSA_{kt}^C based on actual contributions made by those who sign up for the HDHP. This model yields a family-specific prediction of incremental HSA_k^C, denoted $\widehat{HSA_k^C}$, which is inserted into the model such that $P_{k,HDHP} = HSA_k^S + \tau_k \widehat{HSA_k^C}$. Appendix E discusses this model in detail.

Given this setup, we follow the literature and assume that families have constant absolute risk aversion (CARA) preferences implying that, for a given ex post consumption level x:[34]

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

Here, γ_k is a family-specific risk preference parameter that is known to the family but unobserved to the econometrician. We model this as a function of employee demographics X_k^A. As γ increases, the curvature of u increases and the decision maker is more risk averse. The CARA specification implies that the level of absolute risk aversion $\frac{-u''(\cdot)}{u'(\cdot)}$, which equals γ, is constant with respect to the level of x (and, thus, W_k).[35]

In our baseline empirical specification a family's overall level of consumption x conditional on a draw s from $F_{kj}(\cdot)$ is:

---

[33]Incremental contributions to the HSA have value equal to $\tau_{kt} HSA_{kt}^C$ 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 explicitly.

[34]In addition to the reasons the literature assumes CARA risk preferences (such as simplicity) it is important for us to use CARA so that our analysis of adding information frictions is an 'apples to apples' comparison to prior work.

[35]The measure for W_k would matter for an alternative model such as constant relative risk aversion (CRRA).

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

Here, ϵ_{kj} is a family-plan specific idiosyncratic preference shock that is assumed to be mean zero in estimation. Subject to this model, families choose the plan j that maximizes U_{kj}.

There are several key assumptions in the baseline model. First, it assumes that families know the distributions of their future health expenditure risk F_{kj} and that this risk conforms to the output of the cost model described later in this section. This presumes that consumers (i) are fully informed about their own health risk and (ii) fully understand the mapping between total health expenses and out-of-pocket expenses in each plan. The first assumption is violated if, for example, families have private information about their health statuses that is not captured in the prior claims data. Given our detailed individual-level claims data, we believe it is unlikely that there are many consumers with substantial private information in our data. Conversely, given potential difficulties in projecting health risk and expenditures, families may have *less* information about these projections than the econometrician. This possibility, along with the possibility that consumers don't fully understand the health plan characteristics that determine out-of-pocket expenditures, is precisely the kind of issue that motivates the upcoming analysis of information frictions. 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 plan.[36]

Finally, the baseline model also assumes that plans are identical (up to mean zero idiosyncratic ϵ) on non-financial characteristics such as provider network and time/hassle costs. The former is factually correct, though the full model reveals that many lack this knowledge when choosing a plan. For time/hassle costs, we expect there to be differences between the two plans given their respective designs, something that the full model estimates bear out.

Baseline Model With Inertia. One important feature of the choice not captured in the baseline model is inertia. In our setting, if consumers take no action at the time of plan choice in 2012, they will be enrolled in the plan they chose previously as a default option. Prior work (e.g. Handel (2013) and Ericson (2012)) illustrates the 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.

We incorporate inertia into 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_{kj} as follows:

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

---

[36]To the extent that we cannot fully account private consumer health information with our detailed medical data, or limited consumer health information with our survey questions, we perform a robustness analysis in Appendix D to illustrate that our primary estimates are robust to small deviations in F_{kj}.

Here, η represents inertia and depends on observed demographic variables X_k^B, which are described in more detail in the estimation section. $\mathbf{1}_{j_t=j_{t-1}}$ is an indicator for whether the plan you choose this year is the same as your incumbent plan. Apart from the inclusion of η the model with inertia is identical to the baseline model.

There are several assumptions in the model of inertia that warrant discussion. First, inertia is modeled as an incremental cost paid conditional on switching plans (following e.g. Handel (2013), Shum (2004) or Dube et al. (2008)). This implies that, on average, for a family to switch at t they must prefer an alternative option by \$$\eta$ more than their default. There are multiple potential underlying micro-foundations for inertia, each of which could correspond to an alternative model.[37] 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 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 by new employee active choices (conditioning on observable heterogeneity). We present a detailed discussion of identification and estimation after we discuss all the models in this section.[38]

Lastly, we note that information frictions could increase the extent of sub-optimal plan enrollment through both lower quality active decisions and increased inertia. For our primary questions, we care about incorporating inertia into the model along with frictions to better identify risk preferences.[39] Additionally, we are interested in understanding the link between inertia and information frictions. We analyze the extent to which information frictions proxy for inertia if inertia is excluded from the model below. In the model where both inertia and frictions are included, the friction estimates could be interpreted, with some caution, as the 'active choice' impact of frictions above and beyond inertia. We discuss these results and related issues further in Section 5 where we show that (i) information frictions are strong proxies for inertia and (ii) the explicit model of inertia does not have a major impact on the implications for risk preferences estimates in the presence of information frictions.

Full Model The baseline models, with or without inertia, resemble the models examined previously in the structural literature on health insurance markets. Our full model builds on this work by allowing for variation in both consumer information and perceived plan time and hassle costs. While the insight that these factors matter for consumer choice is not new, the ability to measure them and incorporate them into a model with risk preferences and health risk empirically

---

[37]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. First, price changes are not signaled in advance and change little during the study period. Second, it is unlikely that 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). Given this, even if some dynamic considerations exist they should not impact preference estimates markedly.

[38]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.

[39]Further, the baseline model with inertia may be a more realistic portrayal of what prior work in the literature would estimate without the additional choice frictions we incorporate.

is the central innovation of this paper. This is made possible in our setting because of the rich individually-linked survey, claims, and choice data.

There are a multitude of potential ways to incorporate measures of information frictions and hassle costs into our empirical choice model. These span the range from structural to reduced form. A fully structural approach would directly link friction measures derived from the survey to parameters from a model of decision-making under uncertainty subject to limited information. A reduced form approach would include these measures as factors that impact plan valuations without linking them directly to the underlying decision model parameters. In our setting, there is an inherent tension between making additional structural assumptions and the extent to which we must rely on the data to represent specific theoretical parameters. For example, if a consumer incorrectly answers a multiple choice question about what the deductible in the HDHP plan is, we could use the information contained in the answer (e.g. how high or low they answer the deductible is) together with some fairly strong assumptions to estimate a parameter governing how this lack of information directly contributes to the uncertainty in out-of-pocket expenditures represented by $F_{kj}(\cdot)$. Alternatively, a reduced form approach would estimate a shift in valuation for the HDHP plan, relative to the PPO plan, for those who are uninformed relative to those who are informed.

To implement our model, we reduce the number of structural assumptions required and incorporate our survey data using a reduced form approach. Using the data from our linked survey, we construct indicator variables for 'informed', 'uninformed' or 'not sure' answers to each information-relevant survey question as well as variables derived from answers to questions about hassle costs and knowledge of own health expenditures. We include these variables as observable measures of consumer information and perceived hassle costs that imply shifts in value for the HDHP relative to the PPO. For each friction, one category (corresponding to 'no friction', e.g., 'informed') is excluded so that the value shift for the HDHP plan is relative to a frictionless consumer for the measure in question. Specifically, each included friction variable, denoted Z_f from vector $\mathbf{Z}$, shifts the money at stake for each plan, x_j, by an amount $\beta_f Z_f$ that is assumed constant across all potential health state realizations s from $F_j(\cdot)$:

$$x_{kj} = W_k - P_{kj} - s + \eta(X_k^B)\mathbf{1}_{j_t=j_{t-1}} + \mathbf{Z}'_k \beta \mathbf{I}_{HDHP} + \epsilon_{kj}$$

Here, $\mathbf{I}_{HDHP}$ is an indicator variable taking on value of one if plan j is the HDHP plan. To illustrate this setup, if variable Z_1 is an indicator variable that equals 1 if a consumer is uninformed about his deductible, then β_1 measures the difference in valuation for the HDHP plan, for an uninformed person, relative to an informed person. The coefficient β_1 is a reduced form measure that represents the implications of an underlying model of choice under uncertainty with limited information, similar to that presented in Section 2.

The full model includes 13 different variables derived from the survey in the vector $\mathbf{Z}$, not including the excluded 'no friction' categories. These measures are:

- **Information about plan financial characteristics [Questions 1-3 in Table 3]:** We measure whether a person has correct information about HDHP plan financial characteristics. We

construct a binary variable equal to 0 if a consumer knows the deductible, coinsurance rate, and out-of-pocket maximum for the HDHP and a value of 1 otherwise (implying they are at least partially uninformed). A second binary variable has value 1 when a consumer answers 'not sure' to any of these financial characteristic questions, and 0 otherwise.[40] We group knowledge of these financial characteristics together into these two variables because, as shown in Section 3, the answers to these questions are quite correlated.

- **Provider Network Knowledge [Question 7 in Table 4]:** Our next measures study consumer information about the providers that can be accessed in network for each of the two plans. The first (second) variable has value 1 if the consumer believes that one can access more providers/services in the PPO (HDHP). The third equals 1 if the consumer answers 'not sure' to the question on relative provider access. The omitted case is correct knowledge that the plans provide equal access.

- **Information on Own Total Expenditures [Question 10 in Table 4]:** Our next measures study whether a person correctly understands their own total health expenditures. We categorize how an individual's answer about what their expenditures were in the prior year compares to their actual expenditures during that year. 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 expenditures. The omitted case is correct knowledge of past expenditures. We use this measure of past expenditure knowledge to proxy for over or underestimation of projected expenditures for the coming year (the relevant choice object).[41]

- **Tax Benefits Knowledge [Question 6 in Table 3]:** We measure whether or not a consumer understands the tax benefits that a Health Savings Account provides (its main advantage). The first variable equals 1 if the person answers this question incorrectly, while the second one equals 1 if the person answers 'not sure.' The omitted case is the one where the person understands the tax benefits of the HSA.

- **Time and Hassle Costs [Questions 8-9 in Table 4]:** Our final set of measures focuses on stated time/hassle costs and the preferences that consumers have for avoiding them. We develop three variables that describe the interaction between expected time spent on plan logistics / administration and stated preferences for avoiding these activities. The first measure, Z_{HC}, equals the midpoint of the multiple choice option chosen for perceived time spent on HDHP logistics and administration. Thus, if family k answered that they expected to spend '6 to 10 hours' on these activities in the HDHP, $Z_{HCk} = 8$. This variable equals 0 if a consumer expects to spend no time on these activities. The second variable equals $Z_{HC} * \mathbf{I}_D$ while the third equals $Z_{HC} * \mathbf{I}_C$. Here, $\mathbf{I}_D = 1$ if someone states that they ' strongly dislike' spending time on plan

---

[40]We include the separate indicator for 'not sure' vs. 'incorrect' because we believe these answers could be indicative of different types of misinformation.

[41]We asked the question about past expenses, rather than projected future expenses, because we believe questions about past expenditures are simpler than those about future projections. In the latter type heterogeneity in understanding the question and understanding probabilities could swamp a direct measure of under or overestimation.

logistics / administration while $\mathbf{I}_C = 1$ if they answer they are 'concerned about but accept' some time spent on these activities. For preferences, the answer 'don't care about' time spent on these activities is omitted, implying that the coefficient on just Z_{HC} represents the implied time and hassle costs for people with those stated preferences. See Section 3 for an extended discussion.

In addition to estimating the full model described here we estimate a series of models that each include only one of the above frictions, implying five such 'incremental' models. The estimates from these models can be compared to those from the baseline models to see how risk preference estimates are impacted by the inclusion of one additional factor.

For a fully informed individual with no perceived hassle costs the full model with all friction measures reduces to the baseline model (with inertia). Identification and estimation of the full model requires a set of assumptions. First, we assume that risk preferences γ are independent of friction measures $\mathbf{Z}$.[42] Intuitively, this could be violated in either direction. If consumers that are more risk averse give more effort to acquire information, γ will be positively correlated with better information. Conversely, if more sophisticated consumers are generally less risk averse, but acquire information more effectively, this correlation will be negative.[43] As described in the estimation section, we do estimate risk aversion as a function of demographics like age, gender, and income, but this only partially reduces the impact of this assumption. If this assumption is violated, the risk aversion parameters for fully informed consumers will be appropriately identified, but they will be biased for consumers with frictions and that bias will be captured in the parameters β. We note that, if our friction measures capture higher risk aversion for uninformed consumers, our results can be interpreted as upper bounds on the choice and welfare implications of those frictions.

The full model also assumes that frictions shift utility by the same amount for all potential realizations of health expenditures from F. This could be violated if, e.g., someone who believes the PPO plan grants access to more providers believes that lack of access in the HDHP will decrease utility specifically in states where he has a bad health shock. To the extent that this decreases utility for the HDHP *conditional* on risk preferences, we believe that this is appropriately captured in the friction effects β. Further, since the distribution of γ is formally identified with respect to frictionless individuals, this will not be impacted by the perceived utility state-contingency coming from a lack of information.

In addition, the model does not capture correlations between health risk and risk preferences. While prior work (e.g., Cohen and Einav (2007) and Einav et al. (2013)) has illustrated this can be important, especially when thinking about questions related to adverse selection, our primary objective is to estimate shifts in the level of risk preferences when additional frictions are incorporated.

---

[42]We note that a model that conditions risk preferences on $\mathbf{Z}$ would be identified but would substantially increase the number of parameters and, thus, the complexity of the model. Including these correlations would best be done in the 'types' model described shortly, where risk preferences could depend on the one-dimensional information index.

[43]Note here that if consumers have concave utility with respect to uncertainty about plan features (e.g. deductible) this is a different type of risk aversion than that measured by γ. We are concerned with estimating γ, which is risk aversion with respect to out-of-pocket medical expenditures, which insurance is inherently intended to address. We are happy to include the impact of risk aversion with respect to plan characteristics in the coefficients β, since, if forced to enroll in a given plan, we believe this should not be a welfare relevant component of utility, since it does not actually impact marginal utility in good and bad states of the world.

Therefore, we are only concerned about this assumption to the extent that it impacts estimates of this level (though there could be some welfare implications if there is such a correlation).

We do not explicitly model correlations between frictions $\mathbf{Z}$ and inertia η. Depending on the underlying model for inertia, the causal link between inertia and frictions could run in either direction: either inertia leads to a lack of search, and, thus, a lack of information, or, conversely, a lack of clear information provision creates an environment of uncertainty that perpetuates the status quo. In our setting, rather than condition η on $\mathbf{Z}$ and substantially increase the number of parameters, we examine the link between frictions and inertia by estimating the full model with and without inertia. In the model where both inertia and frictions are included, the friction estimates could be interpreted, with some caution, as the 'active choice' impact of frictions above and beyond inertia. In the model without inertia but with frictions, friction estimates also generally capture the extent to which frictions can proxy for inertia. [44]

Types Model. The full model examines the impact of each specific friction measure on choices and on risk preference estimates. Given that we use survey data, rather than administrative data, to measure these frictions, there may be some concerns about how to interpret each survey question or how consumers answering the survey interpret each question. While we designed this survey to be as straightforward as possible to alleviate such concerns, in this section develop a 'types' model that maps the set of disaggregated information frictions described in the last section into a one dimensional index that captures the overall level of information in our environment for a given consumer. This analysis should be more robust; even if there is concern about the interpretation of one or two friction measures, the one dimensional index that aggregates these measures should still be approximately representative of each consumer's level of information. We expect the change in risk preference estimates when information types are included, relative to the baseline model, are similar to those in the full model. The types model is also intrinsically interesting both to understand the distribution of types in the population and to see if there is a strong positive relationship between the overall level of consumer information and choice quality. See e.g., Chetty et al. (2012) for an example of this kind of types analysis in the context of retirement decisions.

We construct our primary type measure as an index that simply adds up the number of information related questions about plan choices that a given consumer gets correct. We use all of the disaggregated frictions in $\mathbf{Z}$ described in the prior section, excluding hassle costs measures, which we still include in the model as a separate friction from the type index.[45] We include each of the financial plan characteristic measures (deductible, coinsurance, out-of-pocket maximum) as

---

[44]This approach works because, in 2011, approximately 85% of consumers enroll in the PPO plan so the β coefficients can also be thought of as generally indexing the non-default plan when inertia is explicitly excluded. When we present our results in Section 5, we show both that information frictions are strong proxies for inertia and that the explicit model of inertia does not have a major impact on the implications for risk preferences estimates in the presence of information frictions.

[45]Time and hassle costs are an important friction to include in all models, but do not have a natural fit into a one dimensional type index with information frictions since they are a distinct type of friction. As discussed in Section 3, consumers could have limited information about the time and hassle costs of a given plan option, the impact of which is incorporated into our estimates.

Figure 4: Histogram of information type index q for the sample of survey respondents.

separate measures, and add two additional measures related to (i) consumer information about the HDHP subsidy and (ii) knowledge about how the HSA compares to the flexible spending account (FSA) (these measures come from questions 4 and 6 in Table 3).

Denote the set of information frictions measures going into the type index as **Z'**. Then the information index q_k is defined:

$$q_k = \Sigma_{Z_f \in \mathbf{Z'}}(1 - Z_f)$$

Here, to be consistent with the notation in the full model, when $Z_f = 1$ this implies a lack of information for a given friction. So, $q_k = 0$ for a completely uninformed consumer and $q_k = 8$ for a completely informed consumer, since we include 8 information related friction measures. Figure 4 plots the distribution of q for the sample of survey respondents. The figure reveals that the distribution of types is skewed towards uninformed, but with substantial heterogeneity and a non-negligible mass of highly informed consumers. In Appendix D we also investigate an alternative information type index q_k' that weights correct answers by the proportion of other consumers who are uninformed (rewarding consumers for degree of difficulty).

The empirical choice model with types is similar to the full model, with the type index replacing the disaggregated frictions. For simplicity, we divide the types q into quartiles ranging from least to most informed.[46] Denote the set of indicator variables, excluding the most informed quartile, as **Q**. Then, the model for money metric utility in each health state s is:

$$x_{kj} = W_k - P_{kj} - s + \eta(X_k^B)\mathbf{1}_{j_t=j_{t-1}} + \mathbf{Q}_k'\beta_Q\mathbf{I}_{HDHP} + \mathbf{Z}_{k,THC}'\beta_{THC}\mathbf{I}_{HDHP} + \epsilon_{kj}$$

Here, the relative utility for the HDHP plan is shifted across all potential health state realizations

---

[46]Given that q_k is discrete, the division into quartiles is not exact but approximates true quartiles.

s by $\mathbf{Q}'\beta_Q$. The set of time and hassle costs measures used in the full model is denoted here as $\mathbf{Z}_{THC}$ and enters the model exactly as before. We note that extensions to the full model, such as correlations between risk preferences and frictions, are more easily captured in the types model, where friction heterogeneity is described in a more parsimonious manner.[47]

Cost Model. The empirical choice framework, for all the specifications presented, takes the distribution of future out-of-pocket expenditures for each family, $F_{kj}(\cdot)$, as given. This section summarizes the empirical model we use to estimate $F_{kj}(\cdot)$, which closely follows the approach used in Handel (2013). Appendix B presents a more detailed description of the model, its estimation algorithm, and its results. The cost model here is intended to estimate the full information, ex ante distribution of out-of-pocket expenditures for each family. Our empirical models accounts for limited information on F by including reduced form dummies for consumers that over or underestimate their own past expenses in survey questions answers.

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. The model has the following primary components:

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.[48] We do this for four distinct types of expenditures: (i) hospital/inpatient (ii) physician office visits (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. 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. We combine individual total expense projections into the family out-of-pocket expense projections used in the choice model, F_{kj}, taking into account family-level plan characteristics.

---

[47] In Appendix D we consider robustness with respect to different representations of the index q, such as sextiles or including the index as constructed without dividing consumers into quartiles.

[48] For example, 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 to fix a time-limited acute condition, such as a broken arm.

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. Both effects are likely to be quite small relative to consumers' total relative valuations of the two plans. Because our cost model combines detailed individual-level prior medical utilization data with sophisticated medical diagnostic software there is less room for private information than in prior work: additional selection based on private information is more unlikely than it would be with only coarse demographics or aggregate health information to measure health risk.[49] To address the question of moral hazard, we perform a robustness analysis in Appendix D that incorporates elasticity estimates from the literature (see e.g., Chandra et al. (2010)) into our cost model. The results show moral hazard impact is small relative to the overall difference in consumers' plan valuations and, therefore, it does not markedly impact our parameter estimates.

Identification. Identification of the empirical parameters in each model is relatively straightforward given the individual-level linked claims data, choice data, and survey data combined with the assumptions for each model. For the baseline model, inertia, information frictions, and time and hassle costs are assumed away and it is assumed that consumer beliefs about future expenditures correspond to the output of the cost model F_{kj}. Subject to F_{kj}, a family's choice in each year identifies a range of feasible risk preferences. 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 similar to how risk preferences are identified throughout the literature (see e.g. Cohen and Einav (2007), Einav et al. (2013), Handel (2013) or Einav et al. (2010a) for a survey).

For the baseline model with inertia, we separately identify inertia from risk preferences by comparing 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. Since this model assumes no additional frictions, it assumes that the two groups are identical on those dimensions. Together with the assumption that the distribution of risk preferences in the population is identical for both groups, conditional on observable heterogeneity, the distinction between new and existing employees identifies inertia separately from risk preferences. Intuitively, choices from active employees identify risk preferences conditional on demographics, and the differences between their choices and those of similar looking existing consumers identify inertia. Table D3 in Appendix D 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 have lower income, and more likely to be single. Importantly, new employees span the ranges of age, gender, and income seen in the full population

---

[49]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 address to some extent with survey data in our full model.

with non-negligible mass, such that estimates of preferences based on observable heterogeneity can credibly be extrapolated from one group to the other.

In theory, for the full model with frictions, inertia is identified in an almost identical manner, but comparisons between new and existing employee choices are now also conditional on observable friction measures. In practice, because the overlap between new employees and survey respondents is small, we estimate inertia using the baseline model with inertia and the full sample of employees at the firm. Then, we include these estimates conditional on observable demographics in the full model. This approach, discussed in detail in the upcoming estimation section, assumes that new employees have the same distribution of information frictions and hassle costs as existing employees. Conditional on inertia, identification of the full model directly follows from the fact that our friction measures are observable variables. The decisions of frictionless consumers identify risk preferences, under the baseline model assumptions on beliefs about the distribution of out-of-pocket expenditures. Then, the coefficients on a given disaggregated friction are identified by comparing the decisions and relative full information plan valuations for consumers with a given friction relative to otherwise identical consumers without that friction (a similar logic holds in the types model). This identification depends on the assumptions that (i) frictions are independent of heterogeneity in risk preferences conditional on observable demographics and (ii) frictions are independent of heterogeneity in inertia conditional on observable demographics.

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:[50]

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

In the primary specifications X_k^A contains employee age, gender, and income. We assume that the family-plan specific error terms ϵ_{kj} 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 of the HDHP relative to that of the PPO.

We assume that the inertia term, $\eta(X_k^B)$ is related linearly to demographics X_k^B:

$$
\eta(X_k^B) = \eta_0 + \eta_1 X_k^B
$$

X_k^B includes income, age, gender, and family insurance coverage tier dummies (corresponding to single, plus one dependent, or two or more dependents).

Our primary estimation sample is the re-weighted survey respondent population described in Section 3 and in Column 3 of Table 1. Due to the limited sample of new employees in the survey

---

[50]We assume that γ is truncated just above zero, at 10^{-15}, though this is generally non-binding.

respondent sample (approximately 100 consumers), we estimate inertia in a first-stage model with the full population of employees and then use these estimates of $\eta(X_k^B)$ as our inertia parameters in the full models. In the first-stage model with the full sample of employees, we estimate the baseline model with inertia where identification of $\eta(X_k^B)$ comes from comparing the choices and relative valuations of existing employees to those of the 2339 new employees in 2011. Denote the estimates of inertia from this model as $\widehat{\eta(X_k^B)}$. We estimate the friction models incorporating $\widehat{\eta(X_k^B)}$ into state-specific money at stake as follows:

$$x_{kj} = W_k - P_{kj} - s + \widehat{\eta(X_k^B)}\mathbf{1}_{j_t=j_{t-1}} + \mathbf{Z}_k'\beta\mathbf{I}_{HDHP} + \epsilon_{kj}$$

It is important to note that the results on how risk preferences estimates (and subsequent welfare implications) change with information frictions are robust to how inertia is incorporated: we show this by comparing the full model results with and without inertia, described in Section 5.

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 coefficients 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 are estimated based on observable heterogeneity, nor the coefficients for information frictions and hassle costs, which are linked directly to the observable survey data. Since the estimation algorithm is similar to a standard approach, we describe the remainder of the details in Appendix C.

5 Results

Table 5 presents the parameters for the baseline specifications with and without inertia, as well as for the first-stage full population model to estimate inertia. In addition to listing the estimated CARA risk preference parameters, the table also provides a simpler interpretation 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 he wins \$1,000 with 50% chance and loses \$$X$ with 50% chance. Thus, if $X = 1,000$, the average consumer is risk neutral, whereas if $X = 0$, the average consumer is infinitely risk averse. In what follows when we refer to 'gamble interpretation' we are referring to the value of X. Bootstrapped standard errors for all parameters are provided in an analogous table, Table D8 in Appendix D.

Column 2 presents the results from the baseline model with the re-weighted sample of survey respondents, our primary sample of interest. The results reveal that the average consumer has what seems like a high degree of risk aversion with $X = \$366.74$ being the amount this consumer would be willing to lose to be just indifferent about accepting the hypothetical gamble.[51] Risk

---

[51]In describing the results, when we refer to 'high' or 'low' risk aversion this is both relative to the other estimates in this paper and relative to other estimates in the literature. It is well known that interpreting specific risk aversion parameters as 'high' or 'low' can be tricky because the economic and welfare implications of those estimates changes with the nature of the specific gamble in question. The welfare results in Section 6 are the true indicators of the

aversion is decreasing in age and income, and is slightly lower for female employees than for male employees: only the age effect is distinct from 0 given the 95% confidence bounds. We estimate substantial unobserved heterogeneity in the risk preferences with σ_γ approximately equal to the average μ_γ.[52] This unobserved heterogeneity becomes much smaller as we move to the full models with information frictions, time/hassle costs, and inertia, suggesting that unobserved heterogeneity in γ may be capturing information that we measure directly in our survey data. The standard deviation of ϵ, the idiosyncratic mean zero heterogeneity is rather small at 149.23, suggesting that unobserved heterogeneity is captured primarily in the risk preference estimates.

Column 1 presents the baseline model with inertia for the full population. The full population model estimates reveal that the average amount of amount of money foregone in plan choice due to inertia is $2,396 with a population standard deviation of $503 based on observable heterogeneity. Figure D1 in Appendix D presents a histogram showing the distribution of estimated inertia in the population as a function of observable heterogeneity while table D4 presents the full set of estimates from the full population inertia model.

The results for the baseline model with inertia for our primary sample, presented in Column 3, illustrate the substantial impact of incorporating inertia on risk preference estimates. The 'gamble interpretation' for the average consumer is $X = 812.61$ (with 95% CI $[733.63, 864.68]$) suggesting that, once inertia is accounted for, the implied level of consumer risk aversion is much lower than that from the baseline model. As in the baseline model, employee age is negatively related to risk aversion, significant at the 95% level, while female and income are also negatively related but with effects small in magnitude and statistically indistinguishable form 0. Notably, σ_γ is now much lower in the model with inertia relative to the baseline model as it is approximately 60% of the much lower average μ_γ. Thus incorporating inertia (with observable heterogeneity) also explains much of the heterogeneity in risk aversion estimated in the baseline model, which foreshadows the effect of including additional friction measures on risk preference heterogeneity.

Table 6 presents the results from our main specifications that include information frictions and time and hassle costs. Standard errors are provided Tables D9, D10, and D11 in Appendix D. The first column repeats the results from the baseline model with inertia, for comparison purposes, while the next four columns present results from the incremental information friction models where we add friction measures one at a time. The final column presents the results from the full model.

Compared to average consumer "gamble interpretation" of $X = 812.61$ for the baseline model with inertia, the mean "gamble interpretation" in the incremental models are $X = 895.35$ for the model with plan financial characteristic frictions, 852.14 with total medical expenditure frictions, 890.42 with provider network / medical access frictions, and 891.16 with time and hassle cost measures included. Except for the model that incorporates total medical expenditure frictions, all incremental models have gamble interpretations for the average consumer that lie outside the 95% confidence interval for that estimated in the baseline model with inertia. Moreover, likelihood ratio

economic consequences of our estimates.

[52]This implies that there is approximately a 30% mass of risk neutral consumers in the model, given that the normal distribution for γ heterogeneity is truncated at 0.

Baseline Models
No Information Frictions

	(1) Full Population Inertia Model	(2) Baseline	(3) Baseline + Inertia
μ_γ - Intercept	$2.01 \cdot 10^{-3}$	$3.21 \cdot 10^{-3}$	$4.15 \cdot 10^{-4}$
μ_γ - Slope, Age	$3.92 \cdot 10^{-7}$	$-3.90 \cdot 10^{-5}$	$-4.66 \cdot 10^{-6}$
μ_γ - Slope, Female	$5.75 \cdot 10^{-5}$	$-1.02 \cdot 10^{-6}$	$-4.58 \cdot 10^{-6}$
μ_γ - Slope, Income	$9.83 \cdot 10^{-7}$	$-1.59 \cdot 10^{-5}$	$-5.57 \cdot 10^{-8}$
Average μ_γ	$2.05 \cdot 10^{-3}$	$1.60 \cdot 10^{-3}$	$2.30 \cdot 10^{-4}$
Std. Dev. μ_γ	$2.47 \cdot 10^{-5}$	$3.09 \cdot 10^{-4}$	$3.64 \cdot 10^{-5}$
Gamble Interpretation of Average μ_γ	305.99	366.74	812.61
σ_γ	$1.70 \cdot 10^{-3}$	$1.79 \cdot 10^{-3}$	$1.57 \cdot 10^{-4}$
σ_ϵ, HDHP	440.29	149.23	5.01
Average Inertia	$2,396*$	-	Int. from (1)
Std. Dev. Inertia	$502*$	-	Int. from (1)

*Detailed estimates of inertia / heterogeneity are in Appendix D.
** Standard errors for all parameters presented in Appendix D.

Table 5: This table presents the structural estimates from (i) the baseline model (ii) the inertial baseline model and (iii) the model that estimates inertia conditional on observable heterogeneity in the full population. The baseline model, described in detail in section 4, estimates health risk, risk preferences, and idiosyncratic preferences while the inertial baseline model integrates the inertia estimates from the full population inertia model. For the full population inertia model we present the population mean and standard deviation of implied inertia, modeled as a switching cost: in Appendix D Table D4 presents all estimates from the full population inertia model.

tests of the incremental models relative to the baseline model with inertia reject the null hypothesis that each of the incremental models is equivalent to that model.

The full model includes all friction measures in one specification: the coefficient estimates on each friction can be interpreted as the average impact of each for choice-relevant valuations. Consumers who believe that the PPO plan has a larger network of medical providers value the HDHP by $2,326 less than someone who correctly knows that these plans grants the same access (significantly different from 0, 95% CI upper bound of -$1,286). Those who underestimate their own total medical expenditures for the past year value the HDHP by $208.30 less than those with correct information while those who overestimate their expenditures prefer the HDHP by $62.98 relative to the fully informed (counter-intuitively). Though the point estimates are wrong signed they are not statistically different from zero. Interestingly, those who answer 'not sure' to this

question value the HDHP by \$688.91 less on average: this may reflect the fact that those who answer 'not sure' have a deep lack of information that causes them to choose the PPO, though there are other potential micro-foundations for this.

Those who answer any of the three main questions on HDHP financial characteristics incorrect actually prefer the HDHP by \$98.04 relative to those who get all of these questions correct, while those who answer 'not sure' to any of these questions have -\$467.48 lower relative average valuations. These effects also have fairly wide 95% CIs that include 0. It is important to note here that while frictions with respect to total medical expenditure knowledge and plan financial characteristic knowledge both have imprecisely estimated coefficients near 0 in the full model, in the incremental models the coefficients for these frictions are negative and large in magnitude, implying a distaste for the HDHP as expected. This suggests that these frictions do imply lower utility for the HDHP plan on their own, but, are overpowered by the other friction measures present in the full model.

Finally, stated time and hassle cost quantities and preferences have a substantial impact on choices. For *each* additional stated hour of time spent on plan billing, administration, and logistics, a consumer with a strong dislike for hassle costs values the HDHP by \$138.70 less. If a consumer "accepts but is concerned about" time and hassle costs, they value the HDHP by \$127.87 less per stated hour. These are relatively precise estimates: the upper bounds on the 95% CIs for these coefficients are -\$79.74 and -\$65.51 respectively. For the median individual in the sample, who expects to incur between 6 and 10 hours of time and hassle costs, this implies (taking the midpoint of 8 hours) a $\$138.70 * 8 = \$1109.60$ drop in utility for the HDHP plan if they state they have a strong dislike for hassle costs. Reassuringly, those who state that they are 'not particularly concerned about' time and hassle costs have a coefficient estimate of \$9.72 less per stated hour which is statistically indistinguishable from 0.

For risk preferences, in the full model we estimate a mean gamble interpretation of $X = 920.47$, lying well above 864.68, the upper bound of the 95% confidence interval for the baseline model with inertia. Moreover, the estimate of σ_γ is substantially reduced relative to the baseline inertial model, suggesting that the heterogeneity estimated in the baseline model proxies for these 'unobservable' frictions . These results demonstrate that, at least in in our setting, having the linked survey data to proxy for information frictions and hassle costs has a economically meaningful and statistically significant impact on estimated risk preferences.

The bottoms rows of the table provides the average and standard deviation of the total effect of frictions on HDHP valuation relative to a perfectly informed consumer. For the incremental models this just implies the average impact of the one friction in question while for the full model this integrates all the friction coefficients multiplied by the appropriate measures. In the full model the average impact on HDHP utility is $-\$1787$ with a standard deviation of 1303.64. This implies that, as expected, on average frictions shift people toward choosing the simpler PPO option. As you move through the incremental models to the full model, the link between the average survey effect and average risk aversion level is apparent: the stronger the mean impact of the frictions the lower the estimated risk aversion.

Table 7 presents the results of the types models with (standard errors presented in Table D12).

Information Frictions & Hassle Costs

Model	(3) Base + Inertia	(4) Plan Design Knowledge	(5) time/hassle Costs	(6) Provider Networks	(7) TME Info	(8) Full Model
Average μ_γ	$2.30 \cdot 10^{-4}$	$1.17 \cdot 10^{-4}$	$1.22 \cdot 10^{-4}$	$1.23 \cdot 10^{-4}$	$1.73 \cdot 10^{-4}$	$8.64 \cdot 10^{-5}$
Std. Dev. μ_γ	$3.64 \cdot 10^{-5}$	$1.75 \cdot 10^{-5}$	$1.73 \cdot 10^{-5}$	$1.86 \cdot 10^{-5}$	$2.47 \cdot 10^{-5}$	$1.39 \cdot 10^{-5}$
Gamble Interpretation of Average μ_γ	812.61	895.35	891.16	890.42	852.14	920.47
σ_γ	$1.57 \cdot 10^{-4}$	$3.88 \cdot 10^{-5}$	$4.79 \cdot 10^{-5}$	$7.35 \cdot 10^{-5}$	$1.02 \cdot 10^{-4}$	$2.19 \cdot 10^{-9}$
σ_ϵ, HDHP	5.01	0.40	30.77	0.52	0.11	17.70
Benefits knowledge:						
Any incorrect	-	-340.79	-	-	-	98.04
Any 'not sure'	-	-777.49	-	-	-	-467.48
Time cost hrs. X prefs:						
Time cost hrs.	-	-	1.65	-	-	-9.72
... X Accept, concerned	-	-	-108.70	-	-	-118.15
... X Dislike	-	-	-137.08	-	-	-128.98
Provider networks:						
HSP network bigger	-	-	-	-1015.50	-	-594.38
PPO network bigger	-	-	-	-2485.84	-	-2362.85
Not sure	-	-	-	-547.85	-	-201.81
TME guess:						
Overestimate	-	-	-	-	-579.50	62.98
Underestimate	-	-	-	-	-674.43	-208.30
Not sure	-	-	-	-	-759.77	-688.91
Average Survey Effect	-	-794.25	-1045.43	-604.21	-391.33	-1787.40
σ Survey Effect	-	328.40	1022.85	757.29	319.83	1303.64
Likelihood Ratio Test Stat vs. (3)	-	64.30	228.63	96.40	44.32	379.54

** Standard errors for all parameters presented in Appendix D.

Table 6: This table presents the results from the primary models with disaggregated 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 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 hassle costs.

The first two columns repeat the baseline model with inertia and the full model for comparison. Reassuringly, the primary types specification estimates, shown in the third column, imply similar implications for frictions as the full model. The average consumer 'gamble interpretation' in the main types model is $X = 930.64$, which is similar to $X = 920.47$ for the full model and, crucially, falls well outside the upper bound of the 95% CI for the baseline model with inertia. The coefficients on time and hassle costs (which are separated from information frictions in the types model) are also similar in both the full model and the types model. The second most informed quartile has an average $1158 dis-utility for the HDHP compared with the most informed consumers (top quartile), with values of $3547 and $6803 for the third and fourth most informed quartiles respectively (for the least informed quartile, very few consumers choose the HDHP). The 95% CIs for each quartile do not overlap with each other. The average effect of frictions on HDHP utility is -$3056 with standard deviation 2299 both somewhat larger that these figures for the full model (resulting from the high negative coefficient on the least informed types). The last column in the table provides a robustness check with a different type index that gives more credit for questions that are difficult to answer for others (described in more detail in Appendix D). The results from this model for risk preferences are similar to those from the primary types model, though the types coefficients are lower in magnitude and σ_ϵ is much higher. Both types models have very high LR test statistics relative to the baseline model with inertia.

The results for the full models just discussed incorporate inertia estimates from the first-stage model. Table D5 in Appendix D examines a model that includes friction measures without the first-stage inertia estimates to (i) examine robustness of the risk preference results with respect to the inertia estimates and (ii) better understand the links between friction measures and inertia.

The full model with disaggregated frictions but no first-stage inertia (column 3) has similar risk preferences estimates to both other full models with inertia, with 'gamble interpretation' for average risk aversion of $X = 914.40$. This illustrates the robustness of our results on risk preferences to the underlying model of inertia; whether we include first-stage estimates or allow frictions to proxy for inertia, the implications for risk preferences, and ultimately welfare, are similar. The average impact of all survey effects is -$3356 (s.d. $1707) approximately $1600 less than that from the model with first-stage inertia estimates. This suggests the our friction measures are good proxies for inertia in our environment. The impact on specific frictions is quite interesting: excluding the first-stage inertia estimates substantially increases the impact of both plan financial knowledge measures and total medical expenditure knowledge measures, while moderately impacting other estimates. This suggests that these two frictions are the most tightly linked to inertia.

Appendix D provides some additional robustness analyses, including an investigation 'placebo' variables derived from the survey and from our administrative data. These analyses verify that risk preference estimates are generally unchanged when including measures that we think should not be predictive of plan choice or utility, such as the number of a building someone works.

Taken together, the results across the estimated models reveal our friction measures both enhance choice predictions and impact risk preference estimates. Additionally, the estimates shed light on which specific frictions may be most important to consumers making health insurance

Aggregated Information Types
& Hassle Costs

Model	(3) None	(8) Full Model	(9) Types Unweighted	(10) Types Weighted
Average μ_γ	$2.30 \cdot 10^{-4}$	$8.64 \cdot 10^{-5}$	$7.45 \cdot 10^{-5}$	$8.95 \cdot 10^{-5}$
Std. Dev. μ_γ	$3.64 \cdot 10^{-5}$	$1.39 \cdot 10^{-5}$	$1.58 \cdot 10^{-5}$	$1.12 \cdot 10^{-5}$
Gamble Interpretation	812.61	920.47	930.64	917.83
σ_γ	$1.57 \cdot 10^{-4}$	$2.19 \cdot 10^{-9}$	$4.91 \cdot 10^{-6}$	$2.10 \cdot 10^{-8}$
σ_ϵ, HDHP	5.01	17.70	0.07	1929.25
Unweighted Information Index*				
Lowest Quartile	-	-	-6803.50	-
Second Quartile	-	-	-3547.10	-
Third Quartile	-	-	-1158.95	-
Weighted Information Index*				
Lowest Quartile	-	-	-	-3655.12
Second Quartile	-	-	-	-1928.53
Third Quartile	-	-	-	-49.46
Time cost hrs. X prefs:				
Time cost hrs.	-	-9.72	-3.09	-33.55
... X Accept, concerned	-	-118.15	-119.99	-101.22
... X Dislike	-	-128.98	-140.73	-107.20
Average Survey Effect	-	-1787.40	-3056.91	-2597.91
SD Survey Effect	-	1303.64	2299.06	1785.42
Likelihood Ratio Test Stat vs. (3)	-	379.54	596.38	707.63

*The omitted category is the fourth quartile, i.e. the most informed consumers.

** Standard errors for all parameters presented in Appendix D.

Table 7: This table presents the results from the 'type' models that aggregate our measures of information frictions into one-dimensional indices that describe the level of information an individual has. Section 4 in the text describes our two type measures 'unweighted' and 'weighted' in more detail. The last two columns in the table present the type models, while the first two columns, for comparison purposes, restate the results from Table 6 of (i) the primary model with inertia but no frictions and (ii) the full model. The overall implications for how risk preference estimates change are similar with the types models and the full disaggregated model, suggesting the type measures are a good representation of underlying heterogeneity in information frictions. Moreover, as expected, the more informed types make 'better' choices on average and are more likely to value the high-deductible plan appropriately.

choices. The change in risk preference estimates as a result of including frictions has a direct impact on welfare analysis and investigation of counterfactual menu design policies, which we now turn to.

6   Policy Analysis: Welfare Impact of Forced HDHP Switch

In this section we study the welfare implications of a counterfactual health insurance menu design where the firm removes the PPO option from the choice set and forces consumers to enroll in the HDHP. While there are a variety of interesting counterfactual plan design scenarios to consider, there are several motivations to study the case of PPO removal. First, the simplicity of this exercise allows us to highlight how the inclusion of data on information frictions and hassle costs can impact welfare analysis and the resulting policy conclusions, without specifying welfare models for information acquisition or consumer inertia. Second, for 2013, the firm that we study decided to change the menu in this exact way, forcing over 40,000 employees to change from the PPO to the HDHP. Third, outside of the context of the large firm we study, other employers face similar menu design choices, and many have chosen to move their employees into high-deductible plans in recent years (see e.g. TowersWatson (2013)). From a broader public policy perspective, regulation of insurance menu design is ubiquitous, with a leading example being the Affordable Care Act (ACA) which legislates the actuarial equivalence values (degree of cost sharing) that private insurance companies can offer to consumers (Kaiser Family Foundation (2011)).

We investigate the welfare impact of forcing consumers into the HDHP plan using the estimates from (i) the baseline model (ii) the baseline model with inertia and (iii) the full model including all disaggregated frictions. For robustness we also examine the full model and no inertia.[53] We compare the welfare implications for these different models to directly illustrate the impact of including data on information frictions and time/hassle costs. Across these models, the primary drivers of welfare, conditional on enrollment, are risk preferences and ex ante distributions of health risk.[54] The key distinction for the full model relative to the baseline models is that the inclusion of information friction measures causes a meaningful change to risk preference estimates *without themselves impacting welfare conditional on enrollment.* Information frictions thus generate a wedge between the demand curve and the 'welfare-relevant' valuation curve which our full model is able to measure (see e.g. Spinnewijn (2012) or Bernheim and Rangel (2009) for an extended discussion of such welfare distinctions). We note that our measures of time/hassle costs could represent either welfare-relevant or non-welfare-relevant micro-foundations, depending on whether these measures reflect true time/hassle costs or misperceptions of them stemming from a lack of information. Our upcoming analysis investigates both of these possibilities.

---

[53]Since the analysis studies movement from a menu offering to a forced choice, inertia is only relevant insofar as it impacts estimates of risk preferences and other welfare-relevant factors when included in the choice model.

[54]Note, in a more general context, differences in provider / treatment coverage would also matter. Here, because the PPO and HDHP are the same on these dimensions, the menu change does not impact welfare analysis on this dimension. The strong consumer choice response to *thinking* the PPO has greater provider access indicates that this is likely an important welfare consideration in general.

Welfare Analysis. We analyze welfare using a certainty equivalent approach that equates the *welfare-relevant* expected utility for each potential health plan option, U_{kj}, with a certain monetary payment Q_{kj}. For the baseline models Q_{kj} solves:

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

The certainty equivalent loss Q_{kj} makes a consumer indifferent between losing Q_{kj} 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_{kj} is a certainty equivalent loss, lower values of Q_{kj} are better from the consumer perspective. For example, $Q = 0$ implies full insurance with no premium while $Q > 0$ implies some cost sharing or premium. For the baseline models, when consumers are forced to join the HDHP, the change in Q coming from the PPO reflects the interaction between risk preferences and the changes in out-of-pocket expenditure risk and plan premium. All choice utility is welfare-relevant in this setup. The mean consumer welfare impact of a policy that forces all consumers into the HDHP is:

$$\Delta CS = \frac{\Sigma_K[Q_{k\hat{j}} - Q_{k,HDHP}]}{K}$$

Here, $Q_{k\hat{j}}$ is the certainty equivalent loss for the plan consumer k chooses to enroll in in the true environment, while $Q_{k,HDHP}$ reflects the certainty equivalent loss when enrolled in the HDHP.

The full model with information frictions and time/hassle costs necessitates a more subtle framework. While the estimation results reveal that the frictions we measure can have a significant impact on choices, we assume that most are not welfare-relevant *conditional on enrollment*. For example, when choosing between the PPO and HDHP, our estimates reveal that 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 (or lack thereof) clearly matters for choice. However, in our counterfactual, when this consumer is forced into the HDHP, the providers are actually the same so welfare from realized provider access should be identical across both plans.

We define $\mathbf{Z}_W$ as subset of frictions that have tangible welfare implications conditional on enrolling in the HDHP and $\mathbf{Z}_{\overline{W}}$ as the complementary subset that do not. For the full model, we construct the welfare-relevant valuations for the different plans by setting the coefficients on $\mathbf{Z}_{\overline{W}}$, $\beta_{\overline{W}}$, equal to 0. These factors thus impact choice utility through $\beta_{\overline{W}}$ but not welfare, driving a wedge between these two quantities. The full model certainty equivalent of enrolling in plan j is:

$$\widehat{U_{kj}}(\gamma_k, \epsilon_{kj}, F_{kj}(\cdot), \mathbf{Z}_W) = u(Q_{kj}) = -\frac{1}{\gamma_k(X_k^A)}e^{-\gamma_k(X_k^A)(W-Q_{kj})}$$

Here, $\widehat{U_{kj}}$ represents the welfare-relevant valuation for consumer k in plan j computed exactly as

U_{kj} in Section 4 but setting non-welfare-relevant coefficients to 0.

While theoretically straightforward, a crucial issue is to determine which frictions are included as welfare-relevant conditional on HDHP enrollment. We believe that, for most of the micro-foundations, there are fairly clear arguments going in one direction or the other. The primary factors in the baseline models, risk preferences and health risk, are clearly welfare-relevant (as they have been throughout the literature).[55] For our friction measures, we assume that all information frictions are non-welfare-relevant conditional on enrollment. We extend the argument given above for knowledge about relative medical care access to the other information frictions to support this assumption. A lack of knowledge about plan financial characteristics or own total medical expenditures impacts choices ex ante, but, conditional on enrollment in the HDHP, this lack of knowledge does not impact the *actual* ex post financial risk faced by consumers in a classical expected utility sense. In essence, we assume that the welfare-relevant utility (or utility conditional on enrollment) is that of a perfectly informed consumer that faces traditional uncertainty with respect to medical expenditures. These assumptions could be violated if, e.g., a lack of information about the deductible or the provider network impacts ex post health care consumption.[56][57]

The most challenging friction to do welfare analysis with is perceived time and hassle costs. If stated time and hassle costs in the HDHP relative to those in the PPO represent true time and hassle costs, then these should be welfare-relevant. Conversely, if the stated measures represent a lack of information about time and hassle costs in the HDHP, and the true values are similar to those in the PPO, the stated measures should not be welfare-relevant: once enrolled in the HDHP a consumer would not actually experience these costs. The analysis in Section 3 suggests that at least part of the high stated HDHP hassle costs are from perceptions rather than true differences. To deal with this issue in the counterfactual analysis, we compute the welfare impact of the forced switch for two scenarios: (i) stated time and hassle costs are full welfare-relevant and (ii) stated time and hassle costs are non-welfare relevant.[58]

Insurance Menu Design: Welfare. For our analysis of the counterfactual policy forcing all consumers to switch to the HDHP, we keep all characteristics of the HDHP constant such that the plan is exactly as in our observed environment. In addition to cost-sharing financial characteristics, this means that we hold the premium constant and don't examine endogenous re-pricing due to a different profile of population health risk. There are two motivations for this. First, we want to highlight the welfare implications of incorporating rich data on information frictions and

---

[55] We also follow convention and assume that idiosyncratic preferences ϵ are welfare-relevant, though these estimates are small in magnitude so this assumption doesn't impact the analysis in any substantial way.

[56] Given generally low estimates in the literature of the price elasticity of medical expenditures, it is likely such ex post responses would not have a major impact on total ex post utility, implying that lack of knowledge on financial dimensions is not likely to markedly impact ex post behavior.

[57] When included in the model, we also assume that knowledge about HSA tax benefits are also non-welfare relevant, though this is a case where it is plausible that this friction could also impact ex post behavior, by causing the consumer to place less money into the HSA. Given the small magnitude of this coefficient, this assumption does not have a major impact on our results.

[58] The case where time and hassle costs are not welfare-relevant could also represent a counterfactual scenario where relative plan hassle costs are reduced to zero.

time/hassle costs, without adding a second dimension of endogenous re-pricing.[59] Second, when the firm actually switched all consumers to the HDHP in 2013, the plan premium remained constant. Our welfare analysis is presented in terms of the consumer welfare change, given the HDHP plan design. While the absolute level of consumer surplus changes with the relative premium difference between the HDHP and PPO, the relative consumer surplus comparisons across the models are equivalent to the relative total surplus comparisons. All of our counterfactual analysis focuses on the primary survey population, in order to focus on the impact of incorporating information frictions and hassle costs into welfare analysis. In what follows, we present the results of the policy change for the 83% of the population that actually chose the PPO, since the welfare change is 0 by construction for those who originally chose the HDHP.

The top half of Table 8 presents the welfare results for the policy that removes the PPO option from the choice set and forces all consumers to enroll in the HDHP (standard errors are presented in Table D13 in Appendix D).[60] It presents the mean and distributional implications for this welfare impact for the baseline models (with and without inertia) and the full model (with and without inertia). In addition, we present the welfare implications of the forced switch under the assumption that all consumers are risk neutral. Since the primary welfare loss from forcing consumers to switch to the HDHP comes from forcing risk averse consumers to bear more risk, this is useful benchmark that represents the maximum welfare gain for switching people to the HDHP given that plan's design.

The welfare implications of including non-welfare-relevant frictions are evident when comparing the mean consumer surplus changes across the models. The baseline model, where consumers are estimated to be more risk averse, predicts a mean welfare loss of $1,237 from the forced switch, with a population standard deviation of $851. The baseline model with inertia predicts a mean loss of $874 with standard deviation $975. The difference demonstrates that controlling for inertia, even without additional frictions measures, is an important component of preference estimation in our setting. The full model with inertia predicts a mean consumer surplus loss of $788 with standard deviation $1021. Finally, if consumers are all risk neutral the mean loss is $726. Since the different consumer surplus estimates across the models come primarily from differences in risk preference estimates, it is interesting to consider the welfare losses under each model relative to the risk neutral case. The full model predicts $62 lower consumer surplus than the risk neutral case, while the baseline model with inertia predicts a $148 relative loss, almost two and half times as large as that from the full model (the baseline model without inertia difference is $511).[61]

Figure 5 provides a graphical representation of the distributional welfare results contained in

---

[59]Further analysis not done here could investigate the interaction between the additional frictions we model and adverse selection, where the frictions would have a direct impact on choices and, subsequently, premiums, if premiums link directly to the health characteristics of those enrolled. See e.g. Handel (2013) for an analysis along these lines.

[60]The welfare estimates presented here for the full models are those that do not include stated time and hassle costs as welfare relevant. See the discussion earlier in this section, the results when these costs are considered welfare relevant in Appendix D, and the descriptive analysis in Section 3 for an extended treatment of this issue.

[61]The median losses from the forced switch are $1,422 and $953 in the baseline model and baseline model with inertia respectively. This loss is $868 for the full model with inertia, $867 for the full model without inertia, and $791 with risk neutral consumers. The ranking of the welfare loss remains the same for all four models across all of the quantiles examined.

Forced HDHP Enrollment
Welfare Analysis

Model	Mean	St. Dev.	25th	Median	75th	95th
Baseline model, no inertia	-1237.61	851.31	-1833.91	-1422.06	-694.33	288.57
Baseline model	-874.46	975.27	-1691.54	-953.03	-238.38	793.72
Full model, no inertia	-788.33	1021.62	-1651.19	-866.79	-109.89	971.53
Full model	-788.94	1021.47	-1653.4	-867.61	-114.01	968.42
Risk neutral	-726.09	1056.82	-1622.6	-791.78	-24.150	1120.93

Moral Hazard Necessary
To Justify Switch

	Elasticity lower bound	Elasticity upper bound
Baseline model, no inertia	0.280	0.407
Baseline model	0.197	0.286
Full model	0.178	0.258
Risk neutral	0.164	0.237

Table 8: 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. The welfare results are presented for each of five different models to illustrate the impact of incorporating inertia, information frictions, and hassle costs on top of a basic model with 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.

Table 8. It presents the welfare results for the 5th, 10th, 25th, 50th, 75th, 90th, and 95th quantiles in the population, for each model. The baseline model curve represents demand in an inertial setting, the baseline model with inertia curve represents demand in an active choice setting, and the full model curve represents welfare-relevant valuation conditional on enrollment. The figure reveals both that there are substantial distributional implications of the forced switch (not surprising given underlying heterogeneity in health risk) and that incorporating our additional friction measures drives a clear empirical wedge between demand and the welfare-relevant valuation of the HDHP relative to the PPO. Additionally, the similarity between the full model results with and without inertia suggests that (i) our friction measures do an excellent job of proxying for inertia when it is excluded and (ii) that our welfare conclusions in the full model are robust to the inclusion of inertia estimates from the administrative data.

Insurance Menu Design: Policy Implications. One motivation for the firm to switch to the HDHP is to incentivize consumers to reduce wasteful medical expenditures. More generally, this is an underlying reason that many large firms cite when moving employees into high-deductible health
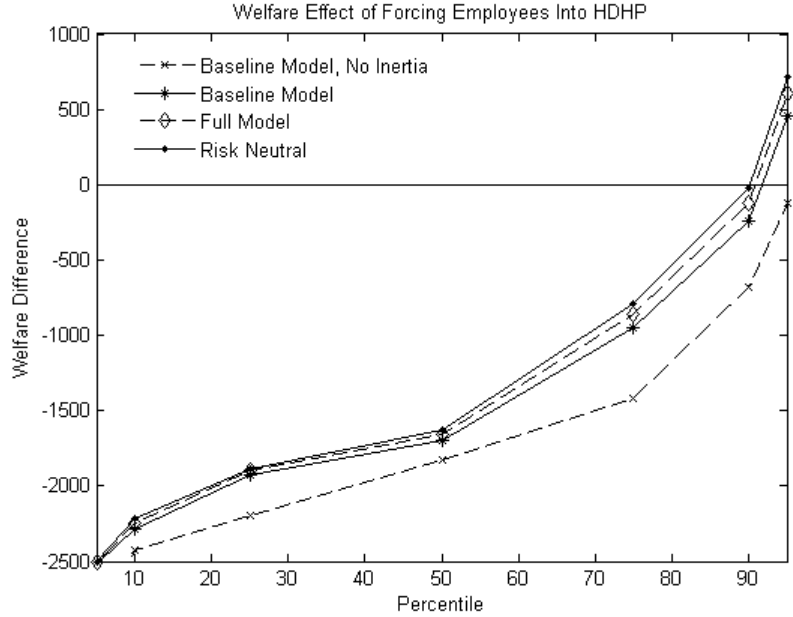
Figure 5: This figure plots quantiles of the welfare impact of the the forced HDHP switch for each of the four models presented in Table 8. The results for both the types model and the full model with no inertia are not included because they heavily overlap with those from the full model presented here: the full model line is a very close representation of the results for each of those models.

plans (see e.g. TowersWatson (2013)).[62] Similarly, some policymakers have called for increased use of HDHPs in an effort to region in health care costs. In order to illustrate the implications of our results from the perspective of a policymaker, or a firm's HR head, we analyze the minimum necessary amount of moral hazard to justify the forced shift to the HDHP. In this calculation, the benefit of reduced wasteful medical expenditures is weighed against the cost of forcing consumers to bear more risk exposure (see e.g. Zeckhauser (1970) for a further discussion).

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 is socially optimal.[63] Since the different models we estimate predict different consumer welfare losses, due primarily to differences in risk preference estimates, 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 becomes larger, a larger price elasticity is necessary to justify that menu redesign.

We calculate the cost savings from reduced medical expenditures in the HDHP due to consumer

---

[62]An additional, off-cited, reason is the desire of large firms to avoid the 'Cadillac Tax' included in the Affordable Care Act (ACA) that taxes plans with high average costs.

[63]Throughout our analysis we have assumed that consumers have a 0 price elasticity for medical utilization. For this exercise, our analysis measures *total cost savings* from a positive elasticity for both the employee and the firm and thus appropriately counts the benefit to the consumer as well as to the firm. As discussed in Section 4 it is unlikely that including a positive utilization elasticity in the choice model would markedly impact the key estimates.

price responsiveness as follows:

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

Here, ΔTC represents the average reduction in total medical expenditures from forcing those enrolled in the PPO to switch into the HDHP. The first fraction measures the mean consumer out-of-pocket price of medical expenditures in the HDHP: this equals 34.9% in our setting and is computed as the average proportion of expenditures paid in the HDHP if consumers had the same total expenditures as in the PPO. μ_{TME_k} denotes mean predicted total medical expenditures for family k for 2012 while $\mu_{F_k,HDHP}$ denotes mean predicted out-of-pocket expenditures for that family in the HDHP that year. The second fraction determines mean total predicted medical expenditures across all families in 2012. We define ξ as the assumed candidate price elasticity for medical expenditures. To simplify our analysis, we assume a homogeneous elasticity in the population (see, e.g., Einav et al. (2013) for estimates of consumer heterogeneity in ξ). Intuitively, the total cost savings from shifting PPO consumers to the HDHP equals the marginal price difference between the those two plans, multiplied by the elasticity ξ to get the proportional reduction in expenditures, and then multiplied by total 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 directly compare ΔTC to the consumer welfare implications from the choice model, we are implicitly assuming that reduced medical expenditures come from reductions in purely wasteful services. In reality, if consumers utilize medical services rationally then they value them by more than their marginal price. 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 ($p_{HDHP} = \frac{\Sigma_K \mu_{F_k,HDHP}}{\Sigma_K \mu_{TME_k,HDHP}}$), rational consumers should value the foregone services at a rate in between 0 and the marginal HDHP price.[64] In this simple model, we can bound the welfare loss to consumers from services foregone below $p_{HDHP} * \Delta TC$. Consequently, we can bound the minimum elasticity necessary to justify the switch to the HDHP between the ξ that equates ΔTC with the change in consumer surplus, ΔCS, and the ξ that equates $(1 - p_{HDHP}) * \Delta TC$ and ΔCS.

The bottom half of Table 8 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 (foregone spending is purely wasteful) while the second column presents the upper bound on this minimum elasticity (foregone spending is valued at consumers' marginal prices). Figure 6 illustrates these calculations in depth. It plots the candidate price elasticity of utilization ξ on the horizontal axis and the resulting total cost savings on the vertical

---

[64]This back of the envelope calculation ignores the fact that the HDHP 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 Cardon and Hendel (2001) for a sophisticated empirical treatment of the demand for medical services. 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: it is plausible, especially given our analysis here, that consumers at times make suboptimal care decisions from a social perspective due to choice and information frictions.
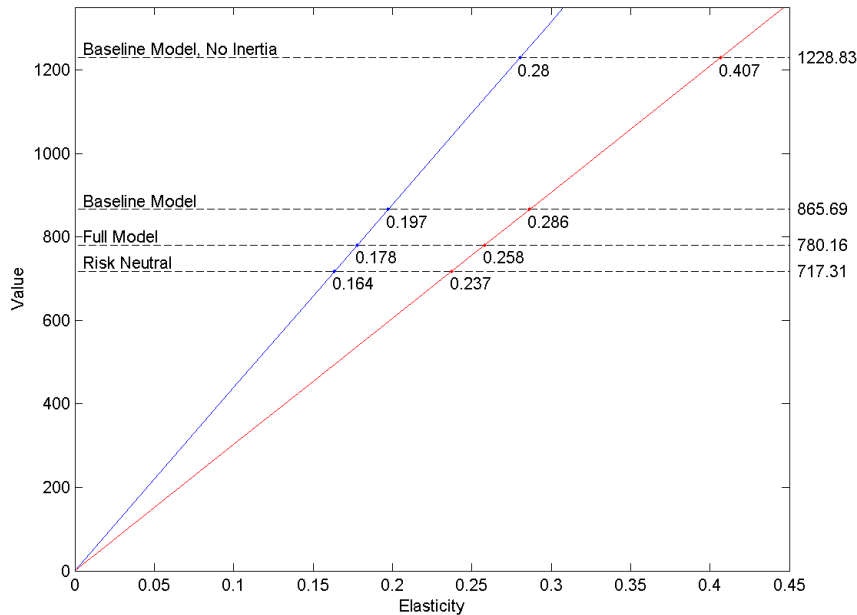
Figure 6: This figure describes the minimum price elasticity necessary to justify the forced switch to the HDHP for each of the choice models studied. The horizontal lines represent the mean consumer surplus change due to increased risk exposure for the different models studied. The upper diagonal line shows the relationship between total cost savings (y-axis) from reduced expenditures in the HDHP for a given price elasticity of utilization ξ (x-axis) when foregone medical expenditures are assumed to be 'purely wasteful.' The second diagonal line shows these costs savings net of the value of foregone under the assumption that that care is valued at the average marginal price paid in the HDHP. These lines imply lower and upper bounds on the minimum price elasticity necessary to justify the menu design change if consumers rationally utilize medical care.

axis, net of the assumed care value foregone by consumers. The figure then compares total cost savings from the policy change to the consumer welfare losses from increased risk exposure implied by each of the different choice models. Since total net cost savings are increasing in ξ, for all ξ greater than some threshold ξ^* the social savings from switching to the HDHP exceed the consumer welfare loss from increased risk exposure under that policy.

The minimum elasticity necessary to justify the menu redesign is 0.280 for the baseline model, when foregone spending is purely wasteful. This elasticity is 0.197 for the inertial baseline model, 0.178 for the full model, and 0.164 under a simple risk neutral model. Thus, as the models incorporate information frictions and hassle costs and the mean consumer welfare loss from risk exposure due to the policy change decreases, this policy becomes more attractive and is justifiable at lower ξ. The upper bound on the minimum elasticity to justify the policy change, when foregone medical care is valued at the marginal price, is 0.407 in the baseline model, 0.286 in the inertial baseline model, 0.258 in the full model, and 0.237 for the risk neutral model.[65]

---

[65]Since these results are based on the case where stated time and hassle costs are assumed to be non-welfare-relevant, or the counterfactual HDHP has the same such costs as the PPO, we note that these minimum elasticities will increase as true HDHP time and hassle costs increase.

While we study a specific large firm and specific menu design counterfactual, the price elasticity thresholds necessary to justify the policy change we consider fall in the range of the elasticity estimates in the literature, such as the oft-cited estimate of 0.18 from the RAND Health Insurance Experiment (see e.g. Chandra et al. (2010) or Newhouse (1993)). While no specific externally valid conclusions should be drawn from these numbers, our analysis illustrates how a simple insurance menu design decision could be directly impacted when the underlying choice model includes detailed measures of information frictions and hassle costs.

7   Conclusion

In this paper we leveraged novel, individually-linked, administrative and survey data to show that both information frictions and perceived hassle costs are important factors for consumer health insurance choices at the large employer we study. We quantified the monetary implications for a variety of specific frictions, and revealed that including these friction measures in an expected utility framework typical of the structural insurance literature has potentially important implications for risk preference estimates. In our setting, omitting the typically unobserved friction measures leads to higher estimates of consumer risk aversion, which in turn directly impacts welfare analysis. In a simple menu design counterfactual analysis designed to highlight the welfare implications of our results, we find that, when we omit our additional friction measures from the model, the consumer welfare loss from risk exposure is approximately double that when these measures are included. While the direction and magnitude of this welfare result are specific to our setting, the analysis illustrates that accounting for these typically unobserved choice frictions can have potentially important implications for both choice and welfare analyses in insurance markets.

Many past studies have noted the potentially important role of information frictions and hassle costs in insurance markets, but few have been able to study even specific frictions in depth, primarily as a result of data limitations. The analysis we perform was made possible by directly linking survey and administrative data at the individual level, and highlights both the additional advantages and potential concerns that leveraging survey data implies. While our analysis attempts to directly address survey elicitation-specific issues like confirmation bias, identification using survey data will always be more open to interpretation that identification using exogenous variation in administrative data. Nevertheless, we argue that integrating survey data with administrative data can produce valuable insights, especially when it is highly unlikely one can obtain rich enough administrative data to answer certain questions, as in our setting.[66]

From a survey design perspective, in our analysis we asked simple questions that consumers could easily understand and took a conservative stance on the exact structural meaning of their answers. An alternative, more ambitious, approach could ask survey questions designed to elicit rich measures of consumer beliefs that could be used directly in a structural framework. This approach would place a greater burden on both question framing and consumer sophistication in

---

[66]This is true broadly for other markets and other empirical questions in economics. See, e.g., Hastings et al. (2013) for an example of linked survey and administrative data in the context of the economics of education.

responding to questions, but could yield a more completely specified structural decision process. An additional extension along these lines could also seek to directly elicit risk preferences as in, e.g., Dohmen et al. (2005). In our analysis, the ability to directly compare our results to the structural health insurance literature is a key benefit of not taking this approach, though integrating such measures could be an interesting complement. In addition, while our analysis is framed in terms of information frictions, there are direct links with the literature on behavioral decision-making that could be further explored (see e.g. Barseghyan et al. (2012) or Abaluck and Gruber (2011)). One could ask survey questions about decision-making directly and use those as inputs into a structural framework. While our framework captures the effects of such decision-making to the extent (i) that the standard expected utility framework provides the "correct" measure of consumer welfare and (ii) our information friction measures are correlated with decision-making biases, clearly there is more room for analysis on this dimension.

Finally, this paper does not focus on the industrial organization implications of our results, which could be interesting to study in future work. In markets where consumers have many potential insurance choices, such as Medicare Part D or the state exchanges proposed in the ACA, the information frictions and relative plan hassle cost differences could be much larger than those we document in our setting with only two primary plan options. While our analysis focused on the demand and welfare implications of measuring such frictions in a simple setting, examining the equilibrium implications of these frictions could be interesting, especially thinking about how firms price to consumers with frictions or the implications of frictions for adverse selection in the marketplace. Additionally, this work has implications for research design: as more exchanges are up and running in 2014 analysts could administer similar surveys in order to better estimate risk preferences, more precisely optimize market design, and answer questions about market equilibrium. Similarly, as large self-insured employers adjust to the ACA and consider whether to move employees towards plans with less risk protection they may be able to perform analyses similar to that here to make more efficient decisions.

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?," Nov. 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.

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. Teitelbaum, "The Nature of Risk Preferences: Evidence from Insurance Choices," November 2012. Cornell Univ. Working Paper.

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

Bertrand, M., E. Duflo, and S. Mullainathan, "How Much Should We Trust Differences-In-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Beshears, J., J. Choi, D. Laibson, and B. Madrian, "How are Preferences Revealed?," *Journal of Public Economics*, 2008, *92*, 1787–1794.

Bundorf, Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 2012, *102* (7), 3214–3248.

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.

Chetty, R., J. Friedman, S. Peterson, T. Nielsen, and T. Olsen, "Active vs. Passive Decisions and Crowdout in Retirement Savings Accounts: Evidence from Denmark," November 2012. NBER Working Paper No. 18565.

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.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. Wagner, "Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey," September 2005. IZA Discussion Paper no. 1730.

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

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

Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 2013, *103* (1), 178–219.

Ericson, K., "Market Design When Firms Interact with Inertial Consumers: Evidence from Medicare Part D," October 2012. NBER Working Paper No. 18089.

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

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

Grubb, M. and M. Osborne, "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," June 2013. M.I.T. Working Paper.

Handel, B., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," March 2013. University of California at Berkeley Working Paper.

Hastings, J., C. Neilson, A. Ramirez, U. Shrestha, and S. Zimmerman, "(Un)informed College Choice: Evidence From Chile," August 2013. Brown University, Working Paper.

Hendren, N., "Private Information and Insurance Rejections," 2013. Harvard Working Paper.

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

Kaiser Family Foundation, "Focus on Health Reform: Summary of the Affordable Care Act," `http://kaiserfamilyfoundation.files.worldpress.com/2011/04/8061-021.pdf` 2011.

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

Kling, Jeffrey, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *Quarterly Journal of Economics*, 2012, *127* (1), 199–235.

Kolstad, J. and A. Kowalski, "Mandate-Based Health Reform and the Labor Market: Evidence from The Massachusetts Reform," 2012. NBER Working Paper 17933.

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.

Manski, C. and S. Lerman, "The Estimation of Choice Probabilities from Choice Based Samples," *Econometrica*, 1977, *45* (8), 1977–1988.

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

Rabin, M., "Psychology and Economics," *Journal of Economic Literature*, 1998, *36*, 11–46.

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.

Solon, G., S. Haider, and J. Wooldridge, "What Are We Weighting For?," February 2013. NBER Working Paper no. 18859.

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

TowersWatson, "Reshaping Health Care: Best Performers Leading the Way," *Towers Watson / NGBH Employer Survey on Purchasing Value in Health Care*, 2013.

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

Zeckhauser, R., "Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives," *Journal of Economic Theory*, 1970, *2* (1), 10–26.

A   Appendix: Survey Instrument

This appendix describes the details of how our survey was administered and provides an exact description of the questions and answer options used in our analyses (described in the text in Tables 3 and 4. The survey was designed in late 2011, in collaboration with the Human Resources (HR) and Communications departments of the employer we study. The team included representatives from a variety of stakeholders within these departments. As described in Section 3, we designed separate surveys for three distinct groups of employees: (i) incumbent HDHP employees (ii) new HDHP enrollees (could have been in PPO before) and (iii) PPO enrollees (there were very few switching back from the HDHP into the PPO from 2011 to 2012). There was substantial overlap in the questions asked to the three groups, although some were irrelevant to a given group and were thus excluded (also, the wording changed to reflect the group in question). Each survey included between 20-25 questions.

The survey was released in early 2012, with electronic invitations sent to 1,500 randomly selected employees from each of the three cohorts above, totaling 4,500 employees. A small group of high-level employees (upper management) were excluded by the HR department as potential survey candidates due to their time constraints. The email was sent from a no-reply address by the employer's insurance provider, and linked to the survey, which was hosted online by this provider. All questions required the employee to choose one or more answers, and never required the employee to fill in their own answers.[67] An example screenshot of two questions from the PPO enrollee survey is given below.

7. If you had signed up for the ▮▮HDHP▮▮ Plan, what would your household's deductible (amount you have to pay for care before the Plan begins to pay for costs) have been this year?
- ○ $0
- ○ $750
- ○ $1500
- ○ $3000
- ○ $3750
- ○ $5000
- ○ G. Not sure

8. If you had enrolled in the ▮▮HDHP▮▮ Plan, what is the rate of coinsurance (% of costs you pay once your deductible is reached) you would pay when visiting an in-network ▮▮▮▮ provider or pharmacy?
- ○ 0%
- ○ 5%
- ○ 10%
- ○ 20%
- ○ 30%
- ○ Not sure

[< Back] [Next >]

All surveys were hosted and completed electronically: respondents were identified when clicking on the link to respond, so that their responses could be linked to the administrative data used in

---

[67]For certain questions that allowed the employee to select one or more answers, an 'Other' option was given. If the employee chose this option, they were prompted to fill in this answer. None of these questions were used in our empirical analysis.

our analysis. As described in the text, we received responses from 579 incumbent HDHP enrollees, 571 new HDHP enrollees, and 511 PPO enrollees for an average response rate of 38%. No financial incentive was given to respond (in the literature, this is quite a high response rate for this kind of survey, given the lack of financial incentive). See Table 1 and the text in Section 3 for detailed comparisons between the full population, survey recipients, and survey respondents on the basis of observable demographics and health risk. The text there discusses respondent selection into the survey, and how it seems minimal on the basis of there observable measures.

We now present the questions and answers used in our analysis, and summarized in the main text in Tables 3 and 4. We present these from the New HDHP enrollee survey and don't present the questions for all three cohorts, since they are very similar to those presented here, with slight wording / framing changes. After delineating these questions, we give a brief discussions of other questions asked but not used in this analysis explicitly. When something is in bold, the true material used was replaced to protect the identity of the firm. For many questions, the order of the answers were shuffled, here was present a specific ordering. The numbering of the questions below corresponds exactly to the numbers for each question used in the main text.

Questions on plan financial characteristics, presented in Table 3 are:

1. What is your household deductible this year in the **HDHP**?

a. $0
b. $750
c. $1,500
d. $3,000
e. $3,750
f. $5,000
g. Not sure

2. In the **HDHP**, what is the rate of coinsurance (% you pay once your deductible is reached) you would need to pay when visiting an in-network **Insurer Name Here** provider or pharmacy?

a. 0%
b. 5%
c. 10%
d. 20%
e. 30%
f. Not sure

3. What is the maximum out-of-pocket you can spend under the **HDHP**, regardless of any funds you or **the firm** may have contributed to your Health Savings Account (HSA)?

a. $0
b. $2,500
c. $5,000
d. $6,250
e. $7,500

f. I don't know

4. How much did **the firm** contribute to your Health Savings Account (HSA) this year, including the Early Adopter Incentive?

a. $0
b. $750
c. $1,500
d. $3,000
e. $3,750
f. $6,250
g. Not sure

5. Which of the following statements is true about the Health Savings Account (HSA)?

a. Funds in the Health Savings Account roll over from year to year
b. If I don't use funds in a given year, they will be lost
c. Not sure

6. Given the tax advantages of a Health Savings Account (HSA), about how much would $1,000 in an HSA be worth in pre-tax dollars in 2012?

a. $700-$999
b. $1,000
c. $1,001-$1,300
d. $1,301-$1,600
e. Greater than $1,600
f. I don't know

The following questions and answers correspond to frictions not related to plan financial characteristics (presented in Table 4 in the main text):

7. How do the medical providers you can use in the **HDHP** in-network compare to those you can use in the **PPO** plan?

a. I can access more providers in the **HDHP**
b. I can access more providers in the **PPO**
c. I can access the same providers under each plan
d. Not sure

8. With any health plan you may spend time choosing medical providers, processing bills, and administering other plan logistics. Approximately how much time do you expect to spend on these activities this year in the **HDHP** plan, assuming a "typical" health year for you and your family?

a. No time at all

b. Less than an hour

c. 1-5 hours

d. 6-10 hours

e. 10-20 hours

f. More than 20 hours

9. Which statement best represents how you feel about spending time managing your **HDHP** plan? (Select One)

a. I understand that I may need to spend time managing my health plan, and I'm not at all concerned about it

b. I accept that I may need to spend time managing my health plan, but I'm concerned with how much time I might have to spend

c. I don't like having to spend time managing my health plan at all, no matter how much time it might be

10. What do you estimate (off the top of your head) is the total cost of the medical care you and your covered dependents consumed (including both what you paid and **the firm** paid) in the last calender year of 2011, i.e. January - December 2011?

a. $0-$500

b. $501-$2,500

c. $2,501-$5,000

d. $5,001-$10,000

e. Greater than $10,000

f. Not sure

11. How confident are you in this estimate (**reference to 10. above**)?

a. Not very confident, or not confident at all

b. Somewhat confident

c. Very confident

12. Based on the total health care needs of you and your dependent(s) in a "typical" year, do you expect to financially benefit from the **HDHP** plan this year (including the value provided by the Health Savings Account and the **firm** contribution)?

a. Yes

b. No

c. Not sure

In addition to these questions, which the analysis focuses on, we ask about 10-15 other ques-

tions covering the following topics:

–Primary reasons for enrolling in HDHP (PPO): consumers can choose several options from list of 7.

–Questions around whether you discussed health plan choice with others, and whether those discussions were informative / influential.

–Questions about consumer learning, including time spent with plan materials provided by Benefits and Communications group and effectiveness of those materials. Also, what plan aspects consumers would like to learn more about.

–Impact of cost-sharing / deductible for medical care utilization. Is utilization impacted by additional cost sharing in HDHP, and, if so, exactly how (list of options)?

Finally, we are currently in the process of running a survey in 2013 that delves more deeply into questions about consumer hassle costs in plan use, consumer medical care utilization, and the mechanisms through which consumers acquire information.

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.[68] 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.[69] 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:

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

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

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

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

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

---

[68]The model is similar to that used in Handel (2013).

[69]In the consumer choice model, this is mostly useful for estimating out-of-pocket expenditures in the HDHP, 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.

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').[70] 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 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 features such as the deductible, co-insurance, co-payments, and out of pocket maximums listed in table A-2. We 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

---

[70]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.

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.

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 total cost 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. Figure B1 outlines the primary components of the cost model pictorially to provide a high-level overview and to ease exposition.

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

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

Once we group individuals into cells for each of the four claims categories, there are two statistical components to estimation. First, we need to generate the continuous marginal distribution of claims for each cell z in claim category d, G_{dz}. To do this, we fit the empirical distribution of claims $G_{I_{dz}}$ to a Weibull distribution with a mass of values at 0. We use the Weibull distribution instead of the 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:

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

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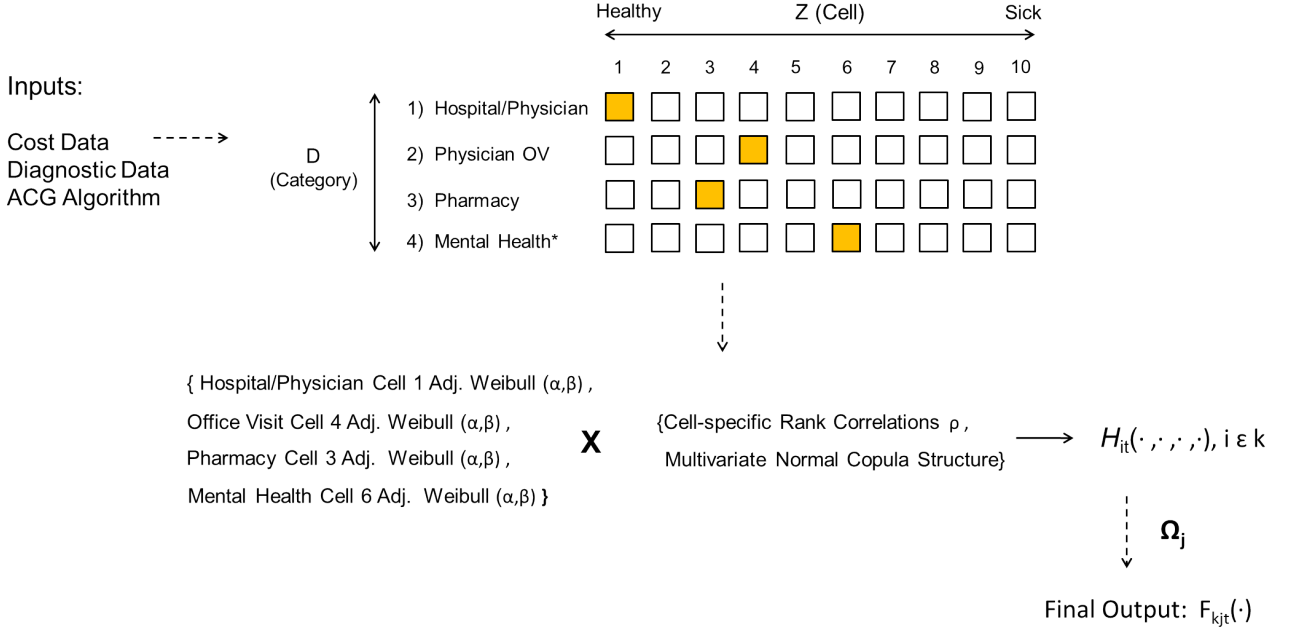
Figure B1: This figure outlines the primary steps of the cost model described in Appendix B. It moves from the initial inputs of cost data, diagnostic data, and the ACG algorithm to the final output F_{kjt} which is the family, plan, time specific distribution of out-of-pocket expenditures that enters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

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.

New Employees. For the first-stage full population model that compares new employees to existing employees to identify the extent of inertia, we need to estimate F_{kj} for new families. Unlike for existing families, we don't observe past medical diagnoses / claims for these families, we just observe these things after they join the firm and after they have made their first health plan choice with the firm. We deal with this issue with a simple process that creates an expected ex ante health status measure. We backdate health status in a Bayesian manner: if a consumer has health status x ex post we construct ex ante health status y as an empirical mixture distribution $f(y|x)$. $f(y|x)$ is estimated empirically and can be thought of as a reverse transition probability (if you are x in period 2, what is the probability you were y in period 1?). Then, for each possible ex ante y, we use the distributions of out-of-pocket expenditures F estimated from the cost model for that type. Thus, the actual distribution used for such employees is described by $\int_{x \in X} f(y|x)F(y)dy$. The actual cost model estimates $F(y)$ do not include new employees and leverages actual claims data for employees who have a past observed year of this data.

C   Appendix: Choice Model Estimation

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

We estimate the choice model using a random coefficients probit simulated maximum likelihood approach similar to that summarized in Train (2009) and to that used in Handel (2013). The simulated maximum likelihood estimation approach has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. We set up a likelihood function to predict the health choices of consumers in 2012. The maximum likelihood estimator selects the parameter values that maximize the similarity between actual choices and choices simulated with the parameters.

First, the estimator simulates Q draws for each family from the distribution of health expenditures output from the cost model, F_k for each family. The estimator also simulates D draws for each family-year from the distribution of the random coefficient γ_k, as well as from the distribution of idiosyncratic preference shocks ϵ_{kj}.

We define θ as the full set of model parameters of interest for the full / primary specification in Section 4:[71]

$$\theta \equiv (\mu_\gamma, \delta, \sigma_\gamma, \sigma_\epsilon, \eta_1, \eta_0, \beta).$$

We denote θ_{dk} as one draw derived from these parameters for each family, including the parameters that are constant across draws (e.g., for observable heterogeneity in γ or η) and those which change with each draw (unobservable heterogeneity in γ and ϵ):[72]

$$\theta_{dk} \equiv (\gamma_k, \epsilon_{kJ}, \eta_k, \beta)$$

Denote θ_{Dk} as the set of all D simulated parameter draws for family k. For each $\theta_{dk} \in \theta Dk$, the estimator uses all Q health draws to compute family-plan-specific expected utilities U_{dkj} following the choice model outlined earlier in section 4. Given these expected utilities for each θ_{dk}, we simulate the probability of choosing plan j^* in each period using a smoothed accept-reject function with the form:

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

This smoothed accept-reject methodology follows that outlined in Train (2009) with some slight modifications to account for the expected utility specification. In theory, conditional on θ_{dk}, we would want to pick the j that maximizes U_{kj} for each family, and then average over D to get final choice probabilities. However, doing this leads to a likelihood function with flat regions, because for small changes in the estimated parameters θ, the discrete choice made does not change. The smoothing function above mimics this process for CARA utility functions: as the smoothing parameter τ becomes large the smoothed Accept-Reject simulator becomes almost identical to the

---

[71]While we discuss estimation for the full model, the logic extends easily to the other specifications estimated in this paper.

[72]Here, we collapse the parameters determining γ_k and η_k into those factors to keep the notation parsimonious.

true accept-reject simulator just described, where the actual utility-maximizing option is chosen with probability one. By choosing τ to be large, an individual will always choose j^* when $\frac{1}{-U_{kj^*}} > \frac{1}{-U_{kj}} \forall j \neq j^*$. The smoothing function is modified from the logit smoothing function in Train (2009) for two reasons: (i) CARA utilities are negative, so the choice should correspond to the utility with the lowest absolute value and (ii) the logit form requires exponentiating the expected utility, which in our case is already the sum of exponential functions (from CARA). This double exponentiating leads to computational issues that our specification overcomes, without any true content change since both models approach the true accept-reject function.

Denote any choice made $\mathbf{j}$ and the set of such choices as $\mathbf{J}$. In the limit as τ grows large the probability of a given $\mathbf{j}$ will either approach 1 or 0 for a given simulated draw d and family k. For all D simulation draws we compute the choice for k with the smoothed accept-reject simulator, denoted $\mathbf{j}_{dk}$. For any set of parameter values θ_{Sk} the probability that the model predicts $\mathbf{j}$ will be chosen by k is:

$$\hat{P}_k^{\mathbf{j}}(\theta, F_{kj}, X_{kt}^A, X_{kt}^B, \mathbf{Z}') = \Sigma_{d \in D} \mathbf{1}[\mathbf{j} = \mathbf{j}_{dk}]$$

Let $\hat{P}_k^{\mathbf{j}}(\theta)$ be shorthand notation for $\hat{P}_k^{\mathbf{j}}(\theta, F_{kj}, X_{kt}^A, X_{kt}^B, \mathbf{Z}')$. Conditional on these probabilities for each k, the simulated log-likelihood value for parameters θ is:

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

Here $d_{k\mathbf{j}}$ is an indicator function equal to one if the actual choice made by family k was $\mathbf{j}$. Then the maximum simulated likelihood estimator (MSLE) is the value of θ in the parameter space Θ that maximizes $SLL(\theta)$. In the results presented in the text, we choose $Q = 50$, $S = 50$, and $\tau = 6$, all values large enough such that the estimated parameters vary little in response to changes.

C.1 Model Implementation and Standard Errors

We implement the estimation algorithm above with the KNITRO constrained optimization package in Matlab. One challenge in non-linear optimization is to ensure that the algorithm finds a global maximum of the likelihood function rather than a local maximum. To this end, we run each model 12 times where, for each model run, the initial parameter values that the optimizer begins its search from are randomly selected from a wide range of reasonable potential values. This allows for robustness with respect to the event that the optimizer finds a local maximum far from the global maximum for a given vector of starting values. We then take the estimates from each of these 12 runs, and select the estimates that have the highest likelihood function value, implying that they are the best estimates (equal to or closest to a global maximum). We ran informal checks to ensure that, for each model, multiple starting values converged to very similar parameters similar to those with the highest likelihood function value, to ensure that we were obtaining robust results.

We compute the standard errors, provided in Appendix D, with a block bootstrap method. This methodology is simple though computationally intensive. First, we construct 50 separate samples, each the same size as our estimation sample, composed of consumers randomly drawn, with replacement, from our actual estimation sample. We then run each model, for 8 different starting values, for each of these 50 bootstrapped samples (implying 400 total estimation runs per model). The 8 starting values are drawn randomly from wide ranges centered at the actual parameter estimates. For each model, and each of the 50 bootstrapped samples, we choose the parameter estimates that have the highest likelihood function value across the 8 runs. This is the final estimate for each bootstrapped sample. Finally, we take these 50 final estimates, across the bootstrapped samples, and calculate the 2.5th and 97.5th percentiles for each parameter and

statistic (we actually use the 4th and 96th percentiles given that 50 is a discrete number). Those percentiles are then, respectively, the upper and lower bounds of the 95% confidence intervals presented in Appendix D. See e.g., Bertrand et al. (2004) for an extended discussion of block bootstrap standard errors.

Finally, it is important to note that the 95% confidence intervals presented in Appendix D should really be interpreted as outer bounds on the true 95% intervals, due to computational issues with non-linear optimization. Due to time and computational constraints, we could only run each of the 50 bootstrap sample runs 8 times, instead of 12. In addition, we could not check each of these bootstrapped runs with the same amount of informal checks as for the primary estimates. This implies that, in certain cases, it is possible that one or several of the 50 estimates for each of the bootstrapped samples are not attaining a global maximum. In this case, e.g., it is possible that 45 of the 50 final estimates are attaining global maxima, while 5 are not. As a result, it is possible that the confidence intervals reported are quite wide due to computational uncertainty, even though the 45 runs that attain the global maximum have results that are quite close together. In essence, in cases where computational issues / uncertainty lead to a final estimate for a bootstrapped sample that is not a global maximum, the confidence intervals will look wide (because of these outlier / incorrect final estimates) when most estimates are quite similar. One solution to this issue would be to run each of the models more times (say 12 or 20) for each bootstrapped sample. This would lead to fewer computational concerns, but would take 1.5 to 2.5 times as long, which is substantial since the standard errors for one model take 7-10 days to run.

As a result, the confidence intervals presented should be thought of as outer bounds on the true 95% CIs. This means that for the models where these bounds are tight, the standard error results are conclusive / compelling since the true 95% CI lies in between these already tight bounds. In cases where the CI is very wide, this means that the true 95% CI lies in that wide range, and that we cannot draw meaningful conclusions due to computational uncertainty in all likelihood. Of course, it is possible the true CI is wide, but, in cases where 46 out of 50 bootstrapped parameter estimates are tight and four are outliers (without substantial variations in the underlying samples) this suggests that computational uncertainty is at fault for the wide bounds.

D Appendix: Additional Analysis

This appendix presents results from additional analyses referred to in the main text. It includes (i) some additional descriptive analysis (ii) several robustness checks for the primary model specifications and (iii) standard errors for all model estimates presented in the main text.

Table D1 presents raw correlations between pairs of binary friction variables derived from the survey. Each entry represents the correlation between the variables listed for the relevant row / column pair. For example, the correlation between correctly knowing one's deductible and correctly knowing one's coinsurance rate is 0.35. As discussed in section 3, the high correlations between several of the friction measures for information about plan financial characteristics suggests that a types specification, which we investigate in section 4 might be interesting.

Table D2 presents the raw correlations between all other frictions measures, including, e.g., perceived time and hassle costs and provider network knowledge. The correlations between these measures are lower, suggesting real heterogeneity in the population across these dimensions. As discussed in the text, this mitigates concerns of confirmation bias. In this table, we include an aggregated measure for knowledge of plan financial characteristics, reflecting the substantial correlations in those measures shown in Table D1 (this is also done here to make the exposition clearer and more parsimonious).

Table D3 presents descriptive statistics for all new employees in 2011, and compares that population to the permanent set of existing employees studied in our full population analyses. The comparison between these two groups is especially relevant to identification of the inertia parameter η in models where it is relevant / included (such full population baseline choice model with inertia, the survey respondent analog to that model, and the sequence of models with friction measures that include η. The table shows that new employees are relatively likely to be younger, lower income, and single. However, they do cover the range of demographics on each of these dimensions in large enough quantities to identify inertia conditional on observable heterogeneity, which mitigates any concerns of selection on these characteristics into the new employee sample for the purposes of estimating inertia. Also, for new employees we include projected health risk distributions that backdate their future (ex-post) claims in a Bayesian manner: i.e. if you have health status x ex post we construct ex ante health status y as an empirical mixture distribution $f(y|x)$. Then, for each possible ex ante y, we use the distributions of out-of-pocket expenditures F estimated from the cost model for that type. Thus, the actual distribution used for such employees is described by $\int_{x \in X} f(y|x)F(y)dy$ (see Appendix B for more details).

Table D4 presents the full results for the first-stage model that estimates risk preferences, health risk, and inertia for the full permanent population references in column 1 of Table 1. The main estimates from this model are presented in the text in Table 5 and discussed there: the table here also includes details regarding the inertia parameter values, including those linked to observable heterogeneity. Figure D1 presents a histogram of the inertia parameter η in the population, where it varies from person to person as a function of observable heterogeneity. The impact of inertia is larger for families than for single employees, reflecting the fact that the former have more money at stake in the health insurance decision. These inertia estimates are used as inputs into the primary models with frictions that we estimate, as described in Section 4.

Table D5 studies the role of inertia in the context of information frictions. The first column presents the results from the baseline model without inertia or information frictions. The second column restates the results from the baseline model with inertia, identified by the choices made by new employees vs. existing employees. The third column presents results from the full model *without* inertia included from the first-stage estimates, while column four repeats the results from the full model with inertia. We include a discussion of this table, and the implications of the

results, in the text in Section 5. The main takeaways are that (i) adding inertia to the baseline model substantially changes risk preference estimates and (ii) when imputed inertia is removed from the full model, the the choice friction estimates become much stronger and replace much of the magnitude of inertia (indicating that inertia is closely related to information frictions). This suggests the our friction measures are good proxies for inertia in our environment. The impact on specific frictions is quite interesting: excluding the first-stage inertia estimates substantially increases the impact of both plan financial knowledge measures and total medical expenditure knowledge measures, while moderately impacting other estimates. This suggests that these two frictions are the most tightly linked to inertia. Finally we note that with or without inertia, the full model has similar risk preference estimates that differ from those in the baseline models.

Figure D2 presents a histogram for an alternative one-dimensional type index to that discussed in the main text (the analogous figure for the primary type index is Figure 4 in Section 4). The alternative index gives consumer more credit if they get "hard" questions correct: specifically, it gives a consumer X points for a correct answer to an information-based question if a (1-X) proportion of the total respondents get that question correct. Thus, for getting a question that no one else gets right correct one gets 1 point, while if everyone else gets the question correct you get 0 points. The two indices are similarly skewed towards uninformed consumers while both have some meaningful mass of informed consumer. Table 7 in the main text includes estimates from a model that includes this alternative type index, along with measures of time and hassle costs. As with the primary type model estimates, there is a monotone relationship between level of information as represented by the index score and consumer valuation of the HDHP plan.

Tables D6 and D7 present two sets of "placebo" models designed as robustness checks to verify that our primary conclusions about the impact of including friction measures on risk preferences are not artifacts of the model setup. I.e., we use placebo models to verify that adding variables that *should be meaningless* don't impact risk preference estimates in any systematic way. The results in these two tables support our framework: adding meaningless placebo variables has little to no impact on risk preference estimates. We use three placebo measures: the first is a random number associated with an employee's actual building. This enters as an actual number that can be related to plan valuation directly: if these numbers are not related to health plan choices and valuations, this variable should not impact risk preference estimates. The additional two placebo measures are (i) a high-level measure of the division of the firm the employee works in (5 such divisions for over 50,000 employees) and (ii) a completely random number. Table D6 presents results when each of these placebo variables is added to the baseline model, without friction measures, while Table D7 presents the results when the placebo variables are added to the full model with all friction measures. Relative to the baseline models, the placebo variables have small coefficients, don't markedly impact risk preference estimates, and actually have negative likelihood ratio test statistics values relative to the baseline model suggesting that these variables add no explanatory power (this reflects estimation uncertainty, in theory, this number should only be positive). The same general conclusions hold for the placebo models relative to the full model, though including these extra variables introduces some estimation uncertainty / difficulties that lead to noisy results.

The remainder of the tables in this appendix present bootstrapped standard errors for all models estimated and discussed in the main text. See Appendix C for a detailed discussion of how standard errors are computed. Here, we present 95% confidence intervals using the block bootstrapped method discussed in that Appendix. We note, as discussed there, that these confidence intervals should be interpreted as bounds on the actual 95% confidence intervals due to estimation uncertainty. For our primary estimates, we ran the estimation routine many times and found the best likelihood function values and also verified that other nearby likelihood results provided essentially identical estimates. For the standard errors, due to computational constraints, we were not able to
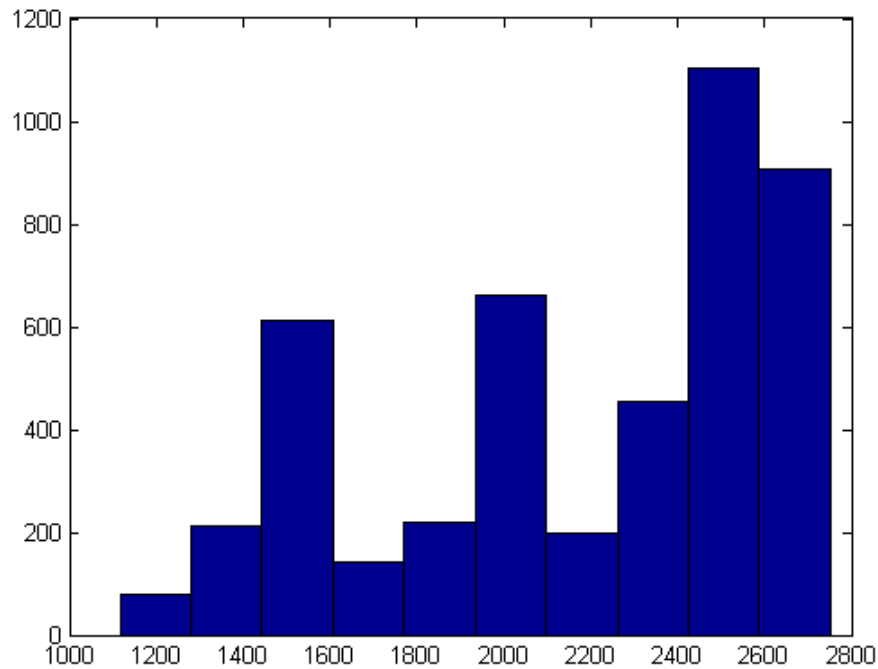
Figure D1: Histogram of Inertia Estimates from full population model, for the full population sample used in that model. Differentiation is based on observable heterogeneity, as seen in Table D4.

run as many estimation runs per sub-sample, leading to additional computational uncertainty. In certain cases, this issue leads to outlier estimation runs (due to finding local maxima rather than global) so it is natural to interpret our intervals as outer bounds on the true CIs in such cases. For many of the specifications, the 95% CI is still quite tight, supporting our main results and allowing meaningful conclusions to be drawn.

Table D8 presents 95% CIs for the set of baseline models, while Tables D9, D10, and D11 presents 95% CIs for all incremental models (with one friction added) and the full model. Finally, Table D12 presents 95% CIs for the two types specifications, and Table D13 presents 95% CIs for the counterfactual simulations run in Section 6. The standard errors and their implications are discussed in the relevant locations in the main text.

Variable	Deductible: Correct	Deductible: Incorrect	Deductible: Not sure	Subsidy: Correct	Subsidy: Incorrect	Subsidy: Not sure	Coinsurance: Correct	Coinsurance: Incorrect	Coinsurance: Not sure
Deductible:									
Correct	–	–	–	0.446	-0.065	-0.364	0.350	0.156	-0.418
Incorrect	–	–	–	-0.091	0.262	-0.124	0.034	0.256	-0.253
Not sure	–	–	–	-0.328	-0.162	0.435	-0.347	-0.357	0.593
Subsidy:									
Correct	0.446	-0.091	-0.328	–	–	–	0.329	0.073	-0.328
Incorrect	-0.065	0.262	-0.162	–	–	–	0.033	0.127	-0.138
Not sure	-0.364	-0.124	0.435	–	–	–	-0.333	-0.169	0.416
Coinsurance:									
Correct	0.350	0.034	-0.347	0.329	0.033	-0.333	–	–	–
Incorrect	0.156	0.256	-0.357	0.073	0.127	-0.169	–	–	–
Not sure	-0.418	-0.253	0.593	-0.328	-0.138	0.416	–	–	–
Out-of-pocket maximum:									
Correct	0.301	0.050	-0.315	0.324	0.051	-0.343	0.324	0.061	-0.313
Incorrect	0.066	0.305	-0.317	0.104	0.129	-0.200	0.121	0.198	-0.272
Not sure	-0.299	-0.295	0.520	-0.349	-0.150	0.445	-0.364	-0.214	0.480
HSA roll over:									
Correct	0.268	0.009	-0.251	0.280	0.058	-0.307	0.169	0.179	-0.293
Incorrect	-0.115	0.080	0.037	-0.127	0.050	0.079	-0.040	-0.051	0.076
Not sure	-0.229	-0.076	0.272	-0.233	-0.110	0.306	-0.171	-0.174	0.290

Variable	OOP maximum: Correct	OOP maximum: Incorrect	OOP maximum: Not sure	HSA roll over: Correct	HSA roll over: Incorrect	HSA roll over: Not sure
Deductible:						
Correct	0.301	0.066	-0.299	0.268	-0.115	-0.229
Incorrect	0.050	0.305	-0.295	0.009	0.080	-0.076
Not sure	-0.315	-0.317	0.520	-0.251	0.037	0.272
Subsidy:						
Correct	0.324	0.104	-0.349	0.280	-0.127	-0.233
Incorrect	0.051	0.129	-0.150	0.058	0.050	-0.110
Not sure	-0.343	-0.200	0.445	-0.307	0.079	0.306
Coinsurance:						
Correct	0.324	0.121	-0.364	0.169	-0.040	-0.171
Incorrect	0.061	0.198	-0.214	0.179	-0.051	-0.174
Not sure	-0.313	-0.272	0.480	-0.293	0.076	0.290
Out-of-pocket maximum:						
Correct	–	–	–	0.197	-0.078	-0.173
Incorrect	–	–	–	0.089	0.058	-0.155
Not sure	–	–	–	-0.233	0.014	0.269
HSA roll over:						
Correct	0.197	0.089	-0.233	–	–	–
Incorrect	-0.078	0.058	0.014	–	–	–
Not sure	-0.173	-0.155	0.269	–	–	–

Table D1: Correlation matrix for responses to information questions on plan financial characteristics. Question responses presented in Table 3.

Variable	Ben. know.: Any incorrect	Ben. know.: Any 'not sure'	Time cost hrs.	Time cost hrs. X Accept	Time cost hrs. X Dislike	Prov. net.: Same	Prov. net.: HSP bigger	Prov. net.: PPO bigger	Prov. net.: Not sure
Benefits knowledge:									
Any incorrect	–	–	0.104	0.070	0.052	0.148	0.101	0.075	-0.243
Any 'not sure'	–	–	-0.131	-0.029	-0.053	-0.177	-0.090	0.002	0.215
Time cost hrs. X prefs:									
Time cost hrs.	0.104	-0.131	–	–	–	0.027	0.072	0.156	-0.164
Time cost hrs. X Accept	0.070	-0.029	–	–	–	0.050	0.070	0.038	-0.107
Time cost hrs. X Dislike	0.052	-0.053	–	–	–	-0.022	0.032	0.132	-0.080
Provider networks:									
Same	0.148	-0.177	0.027	0.050	-0.022	–	–	–	–
HSP network bigger	0.101	-0.090	0.072	0.070	0.032	–	–	–	–
PPO network bigger	0.075	0.002	0.156	0.038	0.132	–	–	–	–
Not sure	-0.243	0.215	-0.164	-0.107	-0.080	–	–	–	–
TME guess:									
Correct	-0.016	-0.152	0.109	0.049	0.036	0.031	0.042	0.008	-0.056
Overestimate	0.032	0.041	0.021	-0.019	0.026	0.061	-0.038	-0.037	-0.017
Underestimate	0.080	0.027	-0.028	0.038	-0.035	0.021	0.027	-0.012	-0.026
Not sure	-0.131	0.148	-0.166	-0.104	-0.048	-0.167	-0.049	0.057	0.149
Tax benefits:									
Understands	0.095	-0.090	0.035	-0.048	0.072	0.135	-0.042	0.001	-0.112
Misunderstands	-0.196	0.259	-0.092	-0.046	-0.029	-0.265	-0.026	0.082	0.217
Not sure	0.130	-0.197	0.068	0.079	-0.022	0.171	0.055	-0.083	-0.140

Variable	TME: Correct	TME: Overestimate	TME: Underestimate	TME guess: Not sure	Tax ben.: Understand	Tax ben.: Misunderstand	Tax ben.: Not sure
Benefits knowledge:							
Any incorrect	-0.016	0.032	0.080	-0.131	0.095	-0.196	0.130
Any 'not sure'	-0.152	0.041	0.027	0.148	-0.090	0.259	-0.197
Time cost hrs. X prefs:							
Time cost hrs.	0.109	0.021	-0.028	-0.166	0.035	-0.092	0.068
Time cost hrs. X Accept	0.049	-0.019	0.038	-0.104	-0.048	-0.046	0.079
Time cost hrs. X Dislike	0.036	0.026	-0.035	-0.048	0.072	-0.029	-0.022
Provider networks:							
Same	0.031	0.061	0.021	-0.167	0.135	-0.265	0.171
HSP network bigger	0.042	-0.038	0.027	-0.049	-0.042	-0.026	0.055
PPO network bigger	0.008	-0.037	-0.012	0.057	0.001	0.082	-0.083
Not sure	-0.056	-0.017	-0.026	0.149	-0.112	0.217	-0.140
TME guess:							
Correct	–	–	–	–	0.023	-0.067	0.051
Overestimate	–	–	–	–	0.009	-0.053	0.047
Underestimate	–	–	–	–	0.007	-0.019	0.015
Not sure	–	–	–	–	-0.060	0.212	-0.171
Tax benefits:							
Understands	0.023	0.009	0.007	-0.060	–	–	–
Misunderstands	-0.067	-0.053	-0.019	0.212	–	–	–
Not sure	0.051	0.047	0.015	-0.171	–	–	–

Table D2: Correlation matrix for responses to plan financial frictions (an aggregated measure) and all other friction measures. Answers to these questions are presented in text in Tables 3 and 4.

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 D3: This table compares employees who are new to the firm in 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.

| **Full Sample** | |
| **Inertia Estimates** | |
	(1)
μ_γ - Intercept	$2.01 \cdot 10^{-3}$
μ_γ - Slope, Age	$3.92 \cdot 10^{-7}$
μ_γ - Slope, Female	$5.75 \cdot 10^{-5}$
μ_γ - Slope, Income	$9.83 \cdot 10^{-7}$
Average μ_γ	$2.05 \cdot 10^{-3}$
Gamble Interpretation of Average μ_γ	305.99
σ_γ	$1.70 \cdot 10^{-3}$
σ_ϵ, HDHP	440.29
Inertia - Intercept	828.16
Inertia - Slope, Age	22.10
Inertia - Slope, Female	-36.69
Inertia - Slope, Income	-32.80
Inertia - Slope, Family size $= 2$	738.69
Inertia - Slope, Family size > 2	1141.11
Average Inertia	2,396
σ Inertia*	502

*The standard deviation reported of inertia in the population is based on observable heterogeneity.

Table D4: This table presents the results the of full population model used to estimate inertia. The identification for inertia in this model comes from comparing the choices made by new employees (with no default option) to those made by existing employees who do have a default option. These inertia estimates are used as inputs into the primary models with frictions, so that the friction impacts are in addition to that linked to inertia. We also estimate a model with friction measures but no inertia in Table D5 which illustrates that, when inertia is not netted out, the friction estimates increase in magnitude, indicating a tight link to inertia, though the change in risk preference estimates are robust to this modeling choice.

Information Frictions and Inertia
Model w/o Explicit Inertia

Model	(1) Baseline	(3) Baseline Inertia	(11) Full Model No Inertia	(8) Full Model Inertia
Average μ_γ	$1.60 \cdot 10^{-3}$	$2.30 \cdot 10^{-4}$	$9.36 \cdot 10^{-5}$	$8.64 \cdot 10^{-5}$
Std. Dev. μ_γ	$3.09 \cdot 10^{-4}$	$3.64 \cdot 10^{-5}$	$1.31 \cdot 10^{-5}$	$1.39 \cdot 10^{-5}$
Gamble Interpretation	366.74	812.61	914.40	920.47
σ_γ	$1.79 \cdot 10^{-3}$	$1.57 \cdot 10^{-4}$	$3.59 \cdot 10^{-9}$	$2.19 \cdot 10^{-9}$
σ_ϵ, HDHP	149.23	5.01	0.05	17.70
Benefits knowledge:				
Any incorrect	-	-	-457.59	98.04
Any 'not sure'	-	-	-1231.69	-467.48
Time cost hrs. X prefs:				
Time cost hrs.	-	-	-58.78	-9.72
... X Accept, concerned	-	-	-99.51	-118.15
... X Dislike	-	-	-95.67	-128.98
Provider networks:				
HSP network bigger	-	-	-777.75	-594.38
PPO network bigger	-	-	-2559.40	-2362.85
Not sure	-	-	-518.66	-201.81
TME guess:				
Overestimate	-	-	-33.18	62.98
Underestimate	-	-	-563.80	-108.30
Not sure	-	-	-1067.37	-688.91
Average Survey Effect	-	-	-3356.28	-1787.40
SD Survey Effect	-	-	1707.11	1303.64
Likelihood Ratio Test Stat vs. (1)	-	1172.75	840.39	1552.29

** Standard errors for all parameters presented in Appendix D.

Table D5: This table studies the role of inertia in the context of information frictions. The first column presents the results from the baseline model without inertia or information frictions. The second column restates the results from the baseline model with inertia, identified by the choices made by new employees vs. existing employees. The third column presents results from the full disaggregated model *without* inertia imputed from the new employee choices, while column four repeats the results from the full model with inertia. The main takeaways are that (i) adding inertia to the baseline model substantially changes risk preference estimates and (ii) when imputed inertia is removed from the full model, the the choice friction estimates become much stronger and replace much of the magnitude of inertia (indicating that inertia is closely related to information frictions). Finally we note that with or without inertia, the full model has similar risk preference estimates that differ from those in the baseline models.

Placebo Tests:
Uninformative Variables
Relative to Baseline

Model	(1) Baseline	(12) Placebo 1 Job Division	(13) Placebo 2 Building ID	(14) Placebo 3 Random Number
Average μ_γ	$1.60 \cdot 10^{-3}$	$1.87 \cdot 10^{-3}$	$1.62 \cdot 10^{-3}$	$1.55 \cdot 10^{-3}$
Std. Dev. μ_γ	$3.09 \cdot 10^{-4}$	$4.25 \cdot 10^{-4}$	$3.53 \cdot 10^{-4}$	$3.40 \cdot 10^{-4}$
Gamble Interpretation	366.74	327.72	363.97	374.56
σ_γ	$1.79 \cdot 10^{-3}$	$2.20 \cdot 10^{-3}$	$1.88 \cdot 10^{-3}$	$1.80 \cdot 10^{-3}$
σ_ϵ, HDHP	149.23	32.61	236.62	152.42
Placebo 1: Job Division*				
Group 1	-	-147.81	-	-
Group 2	-	167.36	-	-
Group 3	-	-132.37	-	-
Group 4	-	5.59	-	-
Group 5	-	-8.24	-	-
Placebo 2: Building ID*				
Group 1	-	-	-178.86	-
Group 2	-	-	-264.49	-
Group 3	-	-	-233.57	-
Placebo 3: Random Number*				
Group 1	-	-	-	-182.49
Group 2	-	-	-	-211.86
Group 3	-	-	-	-93.60
Average Survey Effect	-	-0.61	-142.87	-121.99
SD Survey Effect	-	75.45	109.42	TBD
LR Test Statistic (x) vs. (1)	-	-17.82	-15.54	-21.83

*One category is omitted for each set of placebo variables.

Table D6: This table investigates several 'placebo' models that add what should be meaningless variables to the baseline model. Column 1 repeats the baseline model results, and columns 2-4 describe the placebo model and results, which are discussed in more detail in the text of this appendix. The bottom part of the table investigates hypothesis tests to illustrate that these models are rejected against the models with survey effects. Crucially, the risk preference estimates are unchanged with the addition of placebo variables.

Placebo Tests:
Uninformative Variables
Relative to Full Model

Model	(8) Full Model	(15) Placebo 1 Job Division	(16) Placebo 2 Building ID	(17) Placebo 3 Random Number
Average μ_γ	$8.64 \cdot 10^{-5}$	$3.45 * 10^{-4}$	$1.48 * 10^{-4}$	$1.44 * 10^{-4}$
Std. Dev. μ_γ	$1.39 \cdot 10^{-5}$	TBD	TBD	TBD
Gamble Interpretation	920.47	741.93	871.13	872.05
σ_γ	$2.19 \cdot 10^{-9}$	$2.72 * 10^{-4}$	$6.49 * 10^{-5}$	$5.86 * 10^{-5}$
σ_ϵ, HDHP	17.70	583.89	551.21	620.39
Placebo 1: Job Division*				
Group 1	-	549.52	-	-
Group 2	-	878.39	-	-
Group 3	-	-117.31	-	-
Group 4	-	33.92	-	-
Group 5	-	-85.89	-	-
Placebo 2: Building ID*				
Group 1	-	-	-181.26	-
Group 2	-	-	-252.03	-
Group 3	-	-	464.19	-
Placebo 3: Random Number*				
Group 1	-	-	-	732.81
Group 2	-	-	-	618.73
Group 3	-	-	-	552.14
Average Survey Effect	-1787.40	-1789.73	-2406.10	-2141.23
SD Survey Effect	1303.64	TBD	TBD	TBD
LR Test of vs. (8)	-	-838.55	-756.99	-787.57

*One category is omitted for each set of placebo variables.

Table D7: This table investigates several 'placebo' models that add what should be meaningless variables to the full model with inertia. Column 1 repeats risk preference results from the full model, and columns 2-4 describe the placebo model and results for risk preferences and placebo effects, which are discussed in more detail in the text of this appendix. All friction coefficients are omitted here for brevity, but are available upon request. The bottom part of the table investigates hypothesis tests vs. the full model w/o placebos. The highly negative LR test statistics suggest a lot of estimation uncertainty for these placebo models relative to the full model: in theory these statistics should always be positive though with uncertainty because of complex non-linear optimization this need not be the case in practice.
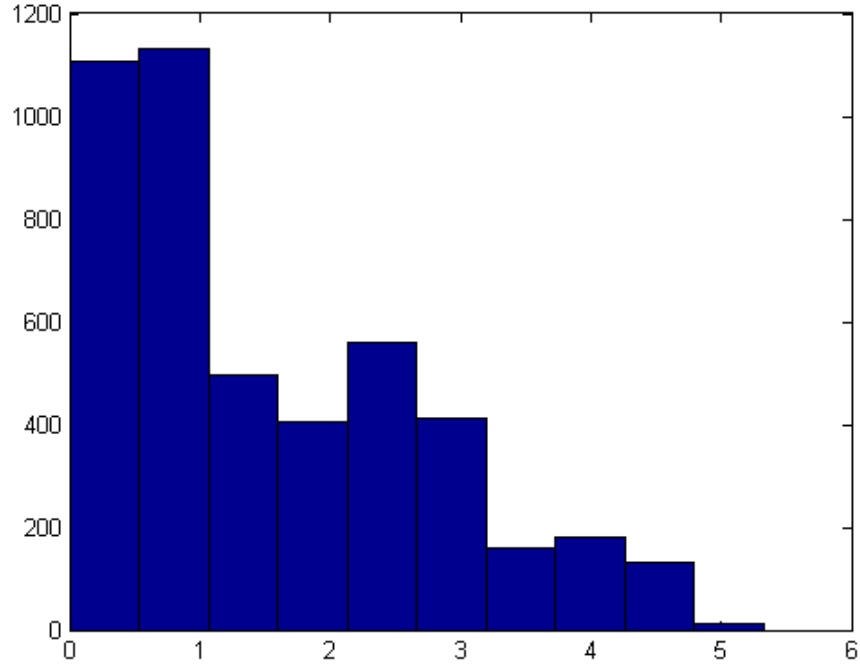
Figure D2: Histogram of weighted information type index q' for the sample of survey respondents.

95% Confidence Intervals
Baseline Models
No Information Frictions

	(2) Baseline	(3) Baseline + Inertia
μ_γ - Intercept	$[3.09 \cdot 10^{-4}, 1.14 \cdot 10^{-2}]$	$[2.73 \cdot 10^{-4}, 6.11 \cdot 10^{-4}]$
μ_γ - Slope, Age	$[-1.38 \cdot 10^{-4}, -2.57 \cdot 10^{-6}]$	$[-7.30 \cdot 10^{-6}, -2.71 \cdot 10^{-6}]$
μ_γ - Slope, Female	$[-1.82 \cdot 10^{-3}, 1.31 \cdot 10^{-3}]$	$[-4.75 \cdot 10^{-5}, 5.06 \cdot 10^{-5}]$
μ_γ - Slope, Income	$[-1.20 \cdot 10^{-4}, 3.87 \cdot 10^{-4}]$	$[-1.31 \cdot 10^{-5}, 1.51 \cdot 10^{-5}]$
Average μ_γ	$[1.69 \cdot 10^{-4}, 7.14 \cdot 10^{-3}]$	$[1.56 \cdot 10^{-4}, 3.64 \cdot 10^{-4}]$
Std. Dev. μ_γ	$[2.64 \cdot 10^{-5}, 1.25 \cdot 10^{-3}]$	$[2.35 \cdot 10^{-7}, 6.38 \cdot 10^{-7}]$
Gamble Interpretation of Average μ_γ	$[97.05, 855.12]$	$[733.63, 864.68]$
σ_γ	$[8.45 \cdot 10^{-5}, 1.04 \cdot 10^{-2}]$	$[6.51 \cdot 10^{-5}, 3.12 \cdot 10^{-4}]$
$\sigma_\epsilon,$	$[0.00, 1513.89]$	$[0.00, 545.13]$

Table D8: This table presents the 95% confidence intervals for the models presented in Table 5 in the main text. The implications of these standard errors are discussed further in Section 5.

95% CIs		
Incremental Models		
Frictions		
	(4)	(5)
Model	Plan Design Knowledge	Time/Hassle Costs
Average μ_γ	$[8.9 \cdot 10^{-5}, 2.2 \cdot 10^{-4}]$	$[1.0 \cdot 10^{-4}, 1.8 \cdot 10^{-4}]$
Std. Dev. μ_γ	$[1.2 \cdot 10^{-5}, 4.1 \cdot 10^{-5}]$	$[1.3 \cdot 10^{-5}, 2.9 \cdot 10^{-5}]$
Gamble Int.	$[821.25, 918.65]$	$[846.37, 907.10]$
of Average μ_γ		
σ_γ	$[0, 1.56 \cdot 10^{-4}]$	$[2.90 \cdot 10^{-5}, 1.07 \cdot 10^{-4}]$
σ_ϵ, HDHP	$[0.00, 63.73]$	$[0.00, 319.45]$
Benefits knowledge:		
Any incorrect	$[-718.23, 163.86]$	–
Any 'not sure'	$[-1191.18, -186.59]$	–
Time costs hrs. X Prefs:		
Time cost hrs.	–	$[-55.61, 108.39]$
... X Concerned	–	$[-206.98, -35.78]$
... X Dislike	–	$[-246.48, -63.58]$
Average Survey Effect	$[-1278.65, -283.61]$	$[-1367.36, -601.63]$
σ Survey Effect	$[150.43, 521.47]$	$[710.47, 1368.95]$

Table D9: This table presents the 95% confidence intervals for the incremental friction model estimates for hassle costs or knowledge of plan financial characteristics, presented in Table 6 in the text. Their implications are discussed further in Section 5.

95% CIs
Incremental Models
Frictions

Model	(6) Provider Networks	(7) TME Info
Average μ_γ	$[1.8 * 10^{-4}, 6.9 * 10^{-3}]$	$[1.8 * 10^{-4}, 8.0 * 10^{-3}]$
Std. Dev. μ_γ	$[2.8 * 10^{-5}, 2.8 * 10^{-3}]$	$[3.3 * 10^{-5}, 2.5 * 10^{-3}]$
Gamble Int. of Average μ_γ	$[100.97, 849.64]$	$[88.13, 845.48]$
σ_γ	$[9.16 * 10^{-5}, 1.12 * 10^{-2}]$	$[1.14 * 10^{-4}, 1.02 * 10^{-2}]$
σ_ϵ, HDHP	$[0, 2953]$	$[23.8, 4226.6]$
Provider networks:		
HSP network bigger	$[-1188.24, -82.59]$	–
PPO network bigger	$[-2865.64, -969.25]$	–
Not sure	$[-1232.05, 887.46]$	–
TME guess:		
Overestimate	–	$[-978, 579]$
Underestimate	–	$[-1082, -63]$
Not sure	–	$[-926, 33]$
Average Survey Effect	$[-928.99, 234.27]$	$[-566.28, 39.48]$
σ Survey Effect	$[450.05, 1052.35]$	$[144.40, 533.26]$

Table D10: This table presents the 95% confidence intervals for the incremental friction model estimates for provider network knowledge or total medical expenditure knowledge, presented in Table 6 in the text. Their implications are discussed further in Section 5.

95% Confidence Intervals	
Full Model	
Disaggregated	
Model	(8) Full Model
Average μ_γ	$[8.19 \cdot 10^{-5}, 2.23 \cdot 10^{-4}]$
Std. Dev. μ_γ	$[9.41 \cdot 10^{-6}, 4.41 \cdot 10^{-5}]$
Gamble Interpretation of Average μ_γ	$[822.51, 924.23]$
σ_γ	$[5.98 \cdot 10^{-6}, 1.55 \cdot 10^{-4}]$
σ_ϵ, HDHP	$[1.58, 666.04]$
Benefits knowledge:	
Any incorrect	$[-614.70, 377.52]$
Any 'not sure'	$[-1670.66, 127.94]$
Time cost hrs. X prefs:	
Time cost hrs.	$[-90.07, 118.86]$
... X Accept, concerned	$[-282.81, -55.79]$
... X Dislike	$[-293.99, -70.02]$
Provider networks:	
HSP network bigger	$[-1842.45, 562.52]$
PPO network bigger	$[-3957.68, -1286.62]$
Not sure	$[-937.44, 303.21]$
TME guess:	
Overestimate	$[-810.72, 704.28]$
Underestimate	$[-1154.63, 837.19]$
Not sure	$[-1987.28, 320.99]$
Average Survey Effect	$[-2148.63, -906.96]$
σ Survey Effect	$[1264.29, 2329.12]$

Table D11: This table presents the 95% confidence intervals for the full model presented in Tables 6 in the text. Implications of these SEs are discussed further in Section 5.

95% Confidence Intervals
Aggregated Information Types
& Hassle Costs

Model	(9) Types Unweighted	(10) Types Weighted
Average μ_γ	$[6.78 * 10^{-5}, 4.26 * 10^{-3}]$	$[6.65 * 10^{-5}, 9.02 * 10^{-5}]$
Std. Dev. μ_γ	$[7.58 * 10^{-6}, 3.74 * 10^{-4}]$	$[3.28 * 10^{-5}, 2.47 * 10^{-3}]$
Gamble Interpretation	[161.7, 936.6],	[917.25, 937.7]
σ_γ	$[3.46 * 10^{-8}, 1.1 * 10^{-2}]$	$[0, 1.72 * 10^{-5}]$
σ_ϵ, HDHP	[0.1,5146]	[0,529.56]
Unweighted Information Index*		
Lowest Quartile	[-8799,-4642]	-
Second Quartile	[-4578,-2613]	-
Third Quartile	[-1879,-625]	-
Weighted Information Index*		
Lowest Quartile	-	[-5334,-3027]
Second Quartile	-	[-3291,-1538]
Third Quartile	-	[-600,410]
Time cost hrs. X prefs:		
Time cost hrs.	[-594,155]	[-123,95]
... X Accept, concerned	[-347,-12]	[-225,-9]
... X Dislike	[-756,-55]	[-245,-27]
Average Survey Effect	[-11,705,-2166]	[-3501,-1980]
SD Survey Effect	[1948,9377]	[1482,2496]

*The omitted category is the fourth quartile, i.e. the most informed consumers.

Table D12: This table presents the 95% confidence intervals for the one-dimensional information types models presented in Table 7. The two models correspond to two different ways to construct the type index, as discussed in the main text. Implications of these SEs are discussed further in Section 5

Model	Mean Welfare Impact Point Estimate	Mean Welfare Impact 95% CI
Baseline model, no inertia	-1237.61	[-1400.70, -807.29]
Baseline model	-874.46	[-970.04, -807.36]
Full model	-788.94	[-923.46 , -695.02]
Risk neutral	-726.09	NA

Table D13: The table presents the 95% CIs for the mean consumer welfare impact of the menu design counterfactual considered in Section 6. See Table 8 for the primary results / discussion in the text.

E   Appendix: Model for Incremental HSA Contributions

This appendix describes how we model incremental employee contributions to their health savings accounts (HSA). The primary model, described in Section 4, incorporates these estimated incremental HSA contributions as inputs into the fixed value / premium that consumers get with the HDHP plan (consumers who enroll in the PPO cannot enroll in or derive value from an HSA).

The primary reason for why we model HSA contributions, rather than use the exact values contributed by each employee (which we observe), is that we need to model the counterfactual contributions PPO enrollees would make if they enrolled in the HDHP. To this end, we train a model of contribution choice based on 2011 HDHP enrollees' actual contributions, and use this model to predict what PPO enrollees might have contributed, were they to enroll in the HDHP. We do the same for actual HDHP enrollees to maintain consistency.

Figures E1, E2, and E3 present the distributions of actual HDP enrollee contributions in 2011, for single employees, employees with one dependent, and employees with more than one dependent respectively. The figures reveal that the distribution of contributions is quite bimodal: either employees choose to forgo contributions altogether, or contribute near the maximum, with very few in between. We note that, for 2013, when all employees were forced to enroll in the HDHP, approximately 60% of employees make positive incremental HSA contributions, a similar proportion to what we see for HDHP enrollees in 2011.

Given their bimodal nature, we model HSA contributions as a two-stage choice. In the first stage, the employee decides whether or not to contribute. Then, if they do decide to contribute, they choose a non-zero amount, which in our model depends on their observable demographics. To estimate the parameters of this model, we first run a probit regression on the decision of 2011 HDHP enrollees to contribute a non-zero amount to their HSA, based on age, gender, income, and family size. We also include a dummy for whether or not their age is above 55, as employees older than that were allowed to contribute an extra "catch-up" $1000 above the normal contribution maximum. We then take those who actually did contribute a nonzero amount, and run a linear regression on their contribution, based on these same demographics. Since employees in different tiers have different maximum contributions, and different incentives to contribute, we run three separate regressions for each coverage tier (single, with spouse, family). The estimates from this model are presented in Table E1.

Based on these estimates, to simulate contributions when estimating the choice model, we
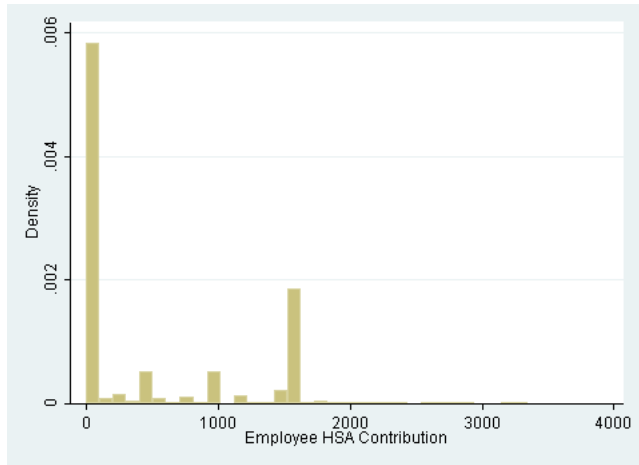
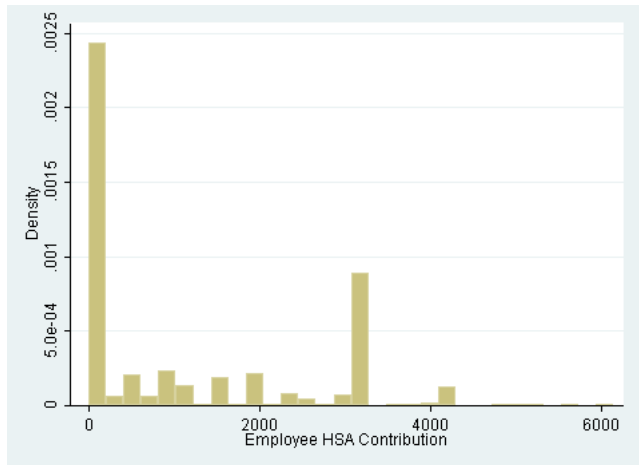Figure E1: Histogram of HSA contributions by single HDHP enrollees in 2011.

Figure E2: Histogram of HSA contributions by employees with one dependent who enroll in the HDHP in 2011.
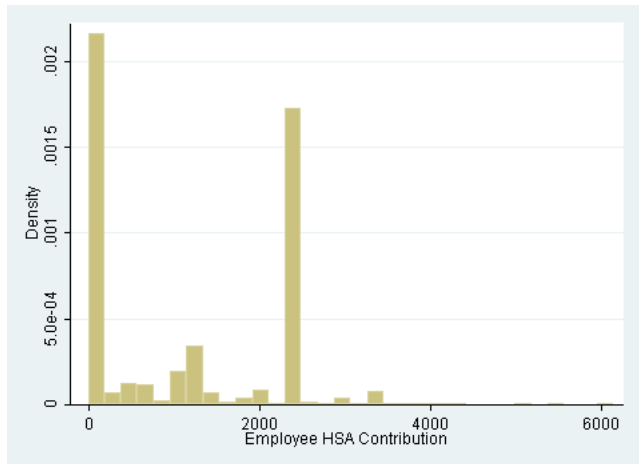

Figure E3: Histogram of HSA contributions by employees with more than one dependent who enroll in the HDHP in 2011.

Model	First Stage	Second Stage Tier 1	Second Stage Tier 2	Second Stage Tier 3
	Probit	OLS	OLS	OLS
Dep. Variable:	HSA > 0	HSA cont.	HSA cont.	HSA cont.
Age	-0.036	11.390	21.823	7.593
Age $\geq$ 55	0.212	439.081	544.984	852.849
Female	0.117	47.305	12.694	19.571
Income	-0.108	68.604	169.556	104.814
Family Tier 2	-0.094	-	-	-
Family Tier 3	-0.062	-	-	-
Intercept	1.279	612.064	849.499	1289.726

Table E1: This table presents the coefficients from the model predicting incremental consumer HSA contributions.

generate a family-specific probability of an HSA contribution based on the first stage. We then draw a Bernoulli random variable with this probability for each family, which determines whether or not they contribute. For those who contribute, their contribution is given by the coefficients coming from second stage associated with their family tier. This output is HSA_k^C is Section 4. Then, the tax benefits from these contributions are obtained by multiplying this contribution by the marginal tax rate τ_k facing the employee, which depends on their observed income level.