

Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets*

Benjamin R. Handel
UC Berkeley and NBER

Jonathan T. Kolstad
UC Berkeley and NBER

Johannes Spinnewijn
London School of Economics and CEPR

January 24, 2018

Abstract

Despite evidence that many consumers in health insurance markets are subject to information frictions, approaches used to evaluate these markets typically assume informed, active consumers. This gap between actual behavior and modeling assumptions has important consequences for positive and normative analysis. We develop a general framework to study insurance market equilibrium in the presence of choice frictions and evaluate key policy interventions, designed to combat adverse selection or to combat poor choices. We identify sufficient relationships between the underlying distributions of consumer (i) costs, (ii) surplus from risk protection and (iii) choice frictions that determine whether friction-reducing policies will be on net welfare increasing, due to improved consumer matching, or welfare reducing, due to increased adverse selection. We show that the impact of insurer risk-adjustment transfers, a supply-side policy designed to combat adverse selection, depends crucially on how effective consumer choices are, and is generally complementary to choice-improving policies. We implement our approach empirically, show how these key sufficient objects can be measured in practice, and illustrate the theoretically-motivated link between these objects and key policy outcomes.

*We thank Glen Weyl for his extensive comments on the paper. We thank Dan Akerberg, Hunt Alcott, Saurabh Bhargava, Stefano DellaVigna, Florian Ederer, Liran Einav, Randall Ellis, Amy Finkelstein, Avi Goldfarb, Josh Gottlieb, Matt Harding, Neale Mahoney, Ariel Pakes, Matthew Rabin, Josh Schwartzstein, Justin Sydnor, Dmitry Taubinsky, and Mike Whinston for their comments. We also thank seminar participants at Arizona State, The Becker Friedman Price Theory Conference, Berkeley-NHH Industrial Organization Conference, Boston University, CEPR, CESifo, Duke, Harvard, KULeuven, Minnesota, Princeton, Tilburg, UCLA, the 2015 Yale Marketing-Industrial Organization Conference, and the ASSA Annual Meetings. We thank Zarek Brot-Goldberg for outstanding research assistance. We thank Microsoft Research for their support of this work.

1 Introduction

A central goal of policy in health insurance markets is to set up an environment whereby firms will offer, and consumers can purchase, efficient insurance products that meet consumer demands for risk protection and health care provision. An important concern in accomplishing this goal is that consumers may be far from fully informed about their health plan choices, and may have difficulties making decisions under limited information [see e.g., Abaluck and Gruber (2011), Ketcham et al. (2012), Kling et al. (2012), Bhargava et al. (2017) and Handel and Kolstad (2015b)]. Despite this observation, the economic models available to evaluate and design common policies in selection markets are by and large based on assumptions of informed, rational consumers [see Chetty and Finkelstein (2013)].

The inability to comprehensively investigate policy impacts in health insurance markets when consumers have meaningful choice frictions is problematic. In practice, a range of policy levers are used to overcome adverse selection, a key impediment to insurance market function [see e.g., Akerlof (1970) or Rothschild and Stiglitz (1976)].¹ Typically, researchers investigating the positive and normative impacts of these policies ignore the potential role of consumer choice frictions. At the same time, regulators also consider and implement policies to reduce choice frictions such as, e.g., information provision, plan recommendations, or smart defaults [see e.g., Handel and Kolstad (2015a)]. These policies are often considered with little or no focus on the impact of adverse selection. Empirical research highlights cases where policies to improve choices can reduce consumer welfare via increased adverse selection [see e.g., Handel (2013)] as well as cases where such policies improve consumer welfare [see e.g. Polyakova (2016)]. However, there has not previously been a systematic investigation of *when* one should expect friction-reducing policies to improve (or decrease) welfare.

In this paper, we develop a general, yet implementable insurance market model that accounts for consumer choice frictions. Our framework allows for the systematic investigation of both policies to combat adverse selection, in the presence of choice frictions, and policies to combat choice frictions in the presence of selection effects. We derive policy-relevant sufficient statistics that identify the key economic tradeoffs and are readily measurable empirically in a wide range of contexts. We demonstrate the applicability of our approach by (i) estimating the relevant primitives in the particular context of employer-provided health insurance and (ii) analyzing the positive and normative impacts of different, oft-considered policies in this setting.

In contrast to prior work that has concentrated on the *mean* value of frictions or inertia [e.g., Handel and Kolstad (2015b), Baicker et al. (2015)], we show that the relative distributions (i.e., mean *and* variance) of three model primitives are crucial for policy and welfare analysis, including (i) the consumers' willingness-to-pay for insurance, (ii) the cost to the insurer and (iii) the impacts of consumer frictions on willingness-to-pay. We map these foundations into demand, cost, and welfare-

¹These regulations include constraints on the types of contracts offered [as in the Affordable Care Act (ACA)], premium subsidies [Einav et al. (2010)], premium risk rating [Bundorf et al. (2012), Handel et al. (2015)], mandates to purchase coverage [e.g., Hackmann et al. (2015)] and risk-adjustment transfers [Cutler and Reber (1998), Handel et al. (2015)].

relevant value curves, building on Einav et al. (2010) and Spinnewijn (2017).² When policies affect the sorting of individuals into insurance, then demand, cost, and welfare curves will change as well and are no longer sufficient to study positive and normative outcomes, as is assumed in most prior work.

We first use our framework to analyze policies that directly reduce consumer choice frictions and identify the following, potentially opposing effects. First, a friction-reducing policy works like a tax when the targeted frictions were pushing consumers at the margin to demand more coverage on average. For that case, the policy worsens under-insurance in an adversely selected market. In addition to this level effect on willingness-to-pay, reducing frictions also affects the sorting of consumers, (i) improving the match between consumers and plans conditional on equilibrium prices and (ii) increasing the equilibrium prices by increasing the correlation between costs and willingness-to-pay. We exploit the tractability of our framework to develop surprisingly simple expressions for the marginal impact of a policy change in terms of means and variances of the demand primitives among the marginal consumers. As the mean and variance of surplus in the population rise relative to those of costs (e.g., due to more heterogeneous preferences), friction-reducing policies become more attractive: the benefits of facilitating better matches between consumers and plans in equilibrium begin to outweigh the costs of increased sorting on costs and subsequent adverse selection. We explore these theoretical properties in a series of simulations designed to highlight these key effects.

In addition to characterizing when friction-reducing policies are ‘good’ or ‘bad’ on their own, we study how these policies interact with the supply-side policy of insurer risk-adjustment transfers. These transfers are designed to reverse adverse selection by compensating insurers who enroll ex ante sicker consumers with transfers from insurers that enroll ex ante healthier consumers. Risk-adjustment transfers are present in many different contexts alongside policies to improve consumer choices (e.g., ACA exchanges, Medicare Part D, Medicare Advantage).³ First, we show that in adversely selected markets increased risk-adjustment improves the impact of friction-reducing policies on welfare, and can shift them from welfare-negative to welfare-positive. Second, we demonstrate that as friction-reducing policies become less attractive (e.g., as the potential for adverse selection increases) effective risk-adjustment plays a much more important role in increasing welfare. These results illustrate the importance of coordinating demand-side interventions with supply-side policies commonly used in insurance markets.

With these insights in hand, we apply our framework to an empirical context where we can measure the distributions of surplus from risk protection, costs, and the impact of frictions on willingness-to-pay. The empirical analysis both highlights the impact the policies we study have

²Our analysis draws a clear distinction between willingness-to-pay and the welfare-relevant valuation once a product is allocated in the spirit of recent work by Baicker et al. (2015) in health care purchasing, Bronnenberg et al. (2014) in generic drug purchasing, Alcott and Taubinsky (2015) in lightbulb purchasing, Rees-Jones and Taubinsky (forthcoming) for tax salience, and Bernheim et al. (2015) in 401(k) allocations. See Dixit and Norman (1978) for a discussion of the distinction between revealed preference and consumer welfare, in the context of advertising.

³See e.g., Cutler and Reber (1998), Brown et al. (2014) or Geruso and McGuire (2016) for discussions of risk-adjustment policies in the literature. See Kaiser Family Foundation (2011) for a discussion of these policies in the context of the ACA.

in one context, and illustrates how our framework can be applied to study similar policy decisions in other contexts. With fairly typical data on individual-level costs and plan designs our framework provides insight into whether friction-reducing policies will be welfare-increasing or welfare-reducing. With more detailed data on key micro-foundations, similar to that in our empirical application, our framework can be used to assess the magnitude of the welfare impact of such policies.

Our empirical analysis builds on the model and estimation in Handel and Kolstad (2015b) using proprietary data on the health plan choices and claims of over 35,000 employees (105,000 employees and dependents) at one large firm, linked at the individual-level to a comprehensive survey designed by the authors to measure the extent of consumers' potentially limited information on many dimensions relevant to health plan choice. Relying on their estimates of risk preferences, health risk, and friction effects on choices we use the data to characterize the non-parametric sample joint distributions of (i) consumer costs, (ii) consumer surplus from risk protection and (iii) the impact of consumer choice frictions on willingness-to-pay. Importantly, we are able to characterize not just the average impact of frictions on willingness-to-pay [the primary focus of Handel and Kolstad (2015b)] but also how they are distributed in the population.

We find that not only is the mean impact of frictions on willingness-to-pay high (mean of \$1,787, pushing consumers towards more generous coverage), but also that the variance in these frictions values is substantial (standard deviation of \$1,304). Expected costs are high, just over \$10,000, as is the variance of costs, implying both high mean and variance of the cost of providing more generous coverage. The mean and variance of estimated surplus from incremental risk protection, however, are both low, reflecting low estimated risk aversion. Given our theoretical results, these foundations suggest that friction-reducing policies on their own will be welfare-reducing: (i) mean friction values are positive and large, so reducing their impact would reduce demand and thus equilibrium coverage, (ii) re-sorting into insurance would be substantial when reducing the heterogeneous impact of frictions and further reduce welfare as the mean and variance of costs are high relative to those of surplus. Thus, informing consumers on their underlying value from insurance will increase the role of cost in decision making, exacerbating adverse selection, without substantially enhancing welfare by allocating people to the plan that gives them more surplus. This also indicates an important role for risk-adjustment transfers as a complement to friction-reducing policies.

These predictions based on our theoretical framework are borne out in our counterfactual analysis. Without any policy interventions, 85% of consumers enroll in more generous coverage with the remaining 15% in just the baseline option. Removing frictions completely, however, leaves only 9% of enrollees in the generous plan, essentially leading to the market fully unraveling, while the surplus of risk-protection is positive for all enrollees. Quantifying the welfare impact, we find that the policy that eliminates frictions reduces the share of first-best surplus achieved to 15%. Risk-adjustment transfers are, however, strongly complementary to friction-reducing policies. When there is no policy in place to reduce frictions, risk adjustment transfers that are 50% (100%) effective increase coverage from 84.6% to 87.1% (88.5%), a positive, but small impact on coverage. However, when the policy to reduce frictions is fully effective, risk adjustment transfers that are

50% (100%) effective increase coverage from 9.1% to 51.6% (63.5%), with similar increases in the percent of first-best surplus achieved. Though the combined policy of fully-reduced frictions and fully-effective risk-adjustment still reduces welfare slightly relative to the status quo, from a distributional standpoint there are fewer consumers leaving substantial sums of money on the table given equilibrium prices.

Our paper proceeds as follows. In Section 2 we present our theoretical framework, characterize market equilibrium and welfare and demonstrate how both are affected by demand-side and supply-side policy interventions. Section 3 describes the data and estimates that we use to empirically implement the model, some descriptive statistics related to consumer heterogeneity on important dimensions and presents our empirical analysis of market equilibrium, friction-reducing policies, and insurer risk-adjustment policies. Section 4 concludes.

2 Theory

Here we develop a stylized model of the insurance market, which can be used to consider available policy options to address adverse selection (e.g., risk adjustment) and information frictions (e.g., consumer choice tools). Focusing on marginal policy changes, we are able to characterize the key trade-offs policy makers are facing and relate them to measurable empirical moments. All proofs are in Appendix A.

2.1 Setup

Our primary model considers a competitive market for one priced insurance plan, following Einav et al. (2010). The plan is offered to all individuals in the market at a uniform price denoted by P . Individuals decide whether to buy the insurance plan or not. An individual i 's willingness-to-pay for the plan is denoted by w_i . Information frictions enter the model as a distortion to individual's willingness-to-pay, following Spinnewijn (2017). The friction, denoted by f_i , results from, e.g., limited information about risks or coverage, or decision-biases at the time of purchase. These frictions are assumed to be exogenous, affecting different individuals differently, potentially induce some to over-estimate the insurance value ($f > 0$) and others to underestimate the insurance value ($f < 0$). The expected cost of providing the coverage depends on the individual's health risk and is denoted by c_i .

We denote the welfare-relevant value of the plan for individual i by $v_i = w_i - f_i$. An individual buys the plan if her willingness-to-pay exceeds the premium, $w_i \geq P$, while her true utility is maximized by buying the plan if and only if $v_i \geq P$. From a welfare perspective, it is efficient for her to buy insurance only if the surplus from risk-protection is positive, $s_i \equiv v_i - c_i \geq 0$.⁴

⁴In an expected utility framework, the value v corresponds to the difference between the certainty equivalent of facing the distribution of total expenses and the certainty equivalent of facing the distribution of out-of-pocket expenses when covered by insurance. The surplus from risk protection will differ for individuals with different risks or preferences. The surplus can also incorporate non-financial plan characteristics. The surplus can in principle be negative due to administrative costs or moral hazard.

Our model thus captures three sources of heterogeneity underlying insurance choices: surplus, cost and frictions. That is, the willingness-to-pay equals

$$w_i = s_i + c_i + f_i.$$

We assume that all demand components are continuously distributed. The additivity of the demand components is not restrictive when we do not impose constraints on the underlying joint distribution.

Our setup could, e.g., reflect a market for supplemental coverage above and beyond a publicly provided government baseline coverage option. The model can also be extended to a market where there are two classes of competitively priced plans (high and low coverage), as studied in Handel et al. (2015) and Weyl and Veiga (2017). We discuss this distinction further in Appendix E. In our context, the comparative statics we study are the same across these distinct setups, though of course actual market outcomes differ. Our setup could also be extended to incorporate issues of moral hazard and imperfect competition, which are empirically relevant in many insurance market contexts.⁵

2.2 Demand, Equilibrium and Welfare

Individuals with different characteristics will sort into insurance depending on the price. The ordering of individuals, and in particular how individuals differ in their characteristics when ordered according to their willingness-to-pay, is key for the analysis. Similarly, any policy intervention that changes the ordering of individuals based on their willingness-to-pay will change the sorting of individuals into insurance and thus affect equilibrium and welfare.

The demand for insurance equals $D(P) = 1 - G(P)$, where G is the cdf of w . We denote the share of buyers by Q . We also introduce the notation $E_P(\cdot) \equiv E(\cdot|w = P)$ and $E_{\geq P}(\cdot) \equiv E(\cdot|w \geq P)$ to denote the expected value among the *marginal* buyers (at the margin between buying insurance or not) and the *infra-marginal* buyers (weakly preferring to buy insurance) respectively.

Our analysis focuses on a competitive environment where the equilibrium price will reflect the expenses made by *all* individuals buying the health plan.⁶ That is, the insurer makes a positive profit as long as the premium P exceeds the average cost of providing insurance to the buyers of insurance at that price, $E_{\geq P}(c)$. Following Einav et al. (2010), we define the competitive price P^c by

$$P^c = E_{\geq P^c}(c). \tag{1}$$

⁵See Mahoney and Weyl (2017) for an analysis of selection markets with imperfect competition (in the absence of frictions), which reverses some typical policy conclusions from competitive selection markets. For moral hazard, see a related discussion in Einav et al. (2010) for the impacts it has in a similar selection markets environment. In our context, including moral hazard would likely have quite limited impacts on positive comparative statics, since willingness-to-pay and costs are typically an order of magnitude larger than the extent of moral hazard. See Brot-Goldberg et al. (2017) for an investigation of moral hazard in our empirical context.

⁶We assume that cost c cannot be observed (or priced) and insurers only compete on prices, taking all other features of the health plan as given. See Veiga and Weyl (2016) and Azevedo and Gottlieb (2017) for an analysis of the plan features provided in equilibrium. Our focus is on consumer frictions and our analysis allows for prices, but no other plan features, to respond to these frictions.

Our focus on this environment is to keep the equilibrium characterization tractable, but several results extend beyond the average-cost pricing we consider.

To evaluate welfare, we consider the total surplus (value net of cost) generated in the insurance market,

$$\mathcal{W}^c = \int_{\tilde{P} \geq P^c} E_{\tilde{P}}(s) dG(\tilde{P}) = [1 - G(P^c)] \times E_{\geq P^c}(s).$$

This criterion assumes that information frictions are not welfare-relevant once a consumer is allocated to a given plan, an assumption we briefly discuss in our empirical context in Section 3. It also ignores distributional consequences of policy interventions, which we briefly consider in the empirical analysis in Section 3.

[FIGURE 1 ABOUT HERE]

Graphical Representation In line with Einav et al. (2010) and Spinnewijn (2017), the market equilibrium and corresponding welfare have a simple graphical representation. We can plot the demand curve $D(P)$ which orders individuals based on their willingness-to-pay and the corresponding marginal cost function $MC(P) = E_P(c)$, average cost function $AC(P) = E_{\geq P}(c)$ and (marginal) value function $V(P) = E_P(v)$. In an adversely selected market, individuals who are more costly to insure have a higher willingness to buy insurance. This causes the cost curve to be downward sloping and the average cost curve to lie above the marginal cost curve, as illustrated in Figure 4. The competitive equilibrium is simply given by the intersection of the demand curve and the average cost curve. To evaluate welfare we need the value of insurance relative to its cost and thus compare the value curve (rather than the demand curve) to the marginal cost curve. Information frictions drive a wedge between the demand curve and the value curve.

2.3 Policy Interventions

We consider the impact of oft-discussed insurance market policies that target (i) improving consumer choices and (ii) reducing adverse selection. To evaluate a policy intervention, it will be useful to decompose its impact into two effects within our framework; a *level effect* effect conditional on the sorting of individuals and a *sorting effect* effect due to the potential re-sorting of individuals. This simple decomposition is useful both at the positive and normative level. A policy can change the equilibrium coverage – either directly or through the re-sorting of individuals based on costs. A policy can change welfare – either through a change in the coverage level Q or by changing sorting into insurance based on surplus, for a given coverage level.⁷

The welfare impact of changing the level of coverage, conditional on the sorting of individuals, is well understood in the literature and simply relates to whether the market is over- or under-insured to start with. In adversely selected markets, average-cost pricing causes the equilibrium price to be inefficiently high and individuals to be *under-insured*. This underlies the analysis of

⁷Since our welfare criterion equals the total surplus, transfers between insurers and insured individuals do not affect welfare. Hence, welfare equals consumer surplus in a competitive equilibrium with zero profits.

price subsidies and mandates in Einav et al. (2010) and Hackmann et al. (2015). However, these studies only considered the pricing inefficiency coming from the supply side. The presence of information frictions may worsen the supply side inefficiency, but can also reduce this inefficiency and potentially reverse the welfare impact of an increase in equilibrium coverage, as argued by Spinnewijn (2017). Frictions can cause individuals to buy coverage even if their valuation is below the price and vice-versa. In particular, if the marginal buyers overestimate the insurance value ($E_P(f) > 0$), this tends to make the equilibrium coverage inefficiently high. The opposite is true if the marginal buyers underestimate the insurance value ($E_P(f) < 0$). The specific welfare impact of different scenarios depends on these offsetting effects, and which dominates.

Proposition 1 *A change in policy x that increases equilibrium coverage $Q(x)$ but maintains the ordering, increases welfare if and only if*

$$[P(x) - E_{P(x)}(c)] - E_{P(x)}(f) \geq 0.$$

The left-hand side equals the marginal surplus at the equilibrium price, $E_{P(x)}(s) = E_{P(x)}(v - c)$, and clearly illustrates the interaction between supply and demand frictions. The marginal surplus equals zero at the *constrained-efficient* price.⁸ From the supply side, insurance companies charge prices that are different from the marginal cost in selection markets, $P(x) \neq E_{P(x)}(c)$. From the demand side, frictions drive a wedge between value and willingness-to-pay, $E_{P(x)}(f) \neq 0$. For example, the under-insurance due to average-cost pricing in an adversely selected market could be fully offset by individuals overestimating the insurance value. But the same friction would further worsen the over-insurance in an advantageously selected market. More generally, it makes clear that policies focused only on the supply side alone may not have their intended effects after accounting for potential demand side frictions. We turn to this later in the context of risk-adjustment transfers.

2.3.1 Information Policies

We first analyze the role of information frictions and how policies that target these frictions depend on the interaction of the demand and supply frictions in selection markets. Improving consumer choices has been a major concern underlying US health care reforms. Regulators and exchange operators have tackled this issues using a number of different policy tools (e.g the provision of information, the regulation and standardization of plan features, the reduction of transaction costs). In our stylized model we consider an information policy that simply reduces the impact of the demand friction f on an individual's willingness to pay. That is,

$$\tilde{w}(\alpha) = w - \alpha \times f$$

with $\alpha \in [0, 1]$ and $\alpha = 1$ capturing the full elimination of demand frictions. An increase in α uniformly reduces the impact of frictions, but this can either increase or decrease an individual's

⁸The *unconstrained* welfare benchmark has individuals sort efficiently and buy insurance if and only if $s \geq 0$.

willingness-to-pay depending on the type of friction affecting her demand.⁹

Level Effect We first consider the level effect of the intervention, conditional on the sorting of consumers. An information policy increases the demand for insurance - just like a subsidy would - when the average friction among the marginal buyers $E_{P(\alpha)}(f)$ is negative. The policy works like a tax if this marginal friction value is positive. Note that even when the average friction value is positive, the marginal friction value can be negative due to the friction-based sorting of individuals. Whether an information policy increases or decreases insurance demand thus crucially depends on the mean and variance of the frictions (in addition to the other primitives affecting the marginal consumers).

Any policy intervention that induces more individuals at the margin to buy insurance decreases the equilibrium price in an adversely selected market (since average cost exceeds marginal cost). This further increases equilibrium coverage. Conditional on the ordering of individuals, an information policy simply scales the impact on quantity of a uniform subsidy, denoted by η^c , depending on the sign and size of the marginal friction value, $E_{P(\alpha)}(f)$.¹⁰ For a uniform friction, this level effect would be the only impact on the market equilibrium and welfare.

Sorting Effect With heterogeneous frictions, an information policy also changes the ordering of individuals' willingness-to-pay. In particular, the policy reduces the willingness for individuals with positive friction values to buy insurance, but increases the willingness for individuals with negative friction values. The impact on the average characteristics of the infra-marginal buyers crucially depends on how those characteristics differ among the marginal buyers with different friction values. For a given share of buyers, the impact of an information policy on $E_{\geq P}(z)$ for any variable z is proportional to the covariance between this variable and the friction value among the marginal buyers, $cov_P(z, f)$.

To illustrate this key result, let us consider the re-sorting on true values first. Among the marginal buyers those with large friction values must have lower true values, while those with low friction values must have higher true values. Hence, a simple selection effect is underlying the re-sorting of individuals; an information policy encourages individuals with high true value to buy insurance and discourages individuals with low true value from buying insurance, as illustrated in Figure 4. The information policy thus necessarily increases the expected true value $E_{\geq P(\alpha)}(v)$ for a given share of buyers. While the re-sorting based on the true insurance value is straightforward, decomposing this sorting effect for costs and surplus is key for positive and normative analysis. The re-sorting based on costs determines the impact on the equilibrium coverage. The re-sorting

⁹ f should be seen as sufficient for any choice policies impacting willingness-to-pay for coverage by αf_i . An extension to the model could consider heterogeneity in α for different policies as well as the underlying heterogeneity in f that we consider here.

¹⁰ As shown in the proof of Proposition, the impact of a uniform subsidy on the equilibrium quantity equals

$$\eta^c \equiv \frac{g(P^c)}{1 - [E_{\geq P^c}(c) - E_{P^c}(c)] \frac{|\varepsilon_D(P^c)|}{P^c}}.$$

based on surplus determines the impact on welfare.

[FIGURE 2 ABOUT HERE]

In an adversely selected market, individuals with higher true valuation have higher expenses, suggesting that the market becomes even more adversely selected when reducing the role of frictions. This would increase the equilibrium price and thus reduce the equilibrium coverage. In general, the impact of re-sorting on equilibrium coverage is captured by the covariance between costs and frictions among the marginal buyers, $cov_{P(\alpha)}(c, f)$. This covariance should be compared to the average friction value among the marginal buyers to assess the impact of the policy intervention on equilibrium coverage.

Regarding welfare, when individuals with higher true valuation have a higher surplus from buying insurance, the average surplus of the individuals buying insurance increases when reducing the frictions (conditional on the share of buyers). The improved matching unambiguously increases welfare, regardless of the nature of competition and whether the equilibrium coverage is efficient or not. In general, the sorting effect is captured by the covariance between the friction value and the surplus among the marginal buyers, $cov_{P(\alpha)}(s, f)$. The total welfare change then depends on this sorting effect in addition to the welfare impact from the change in coverage.

Proposition 2 *An information policy α changes equilibrium coverage in a competitive market by*

$$Q'(\alpha) = -\eta^c \times [E_{P(\alpha)}(f) - cov_{P(\alpha)}(c, f) \frac{|\varepsilon_D(P(\alpha))|}{P(\alpha)}].$$

The corresponding impact on equilibrium welfare equals

$$W'(\alpha) = E_{P(\alpha)}(s) Q'(\alpha) - cov_{P(\alpha)}(s, f) g^{\tilde{w}(\alpha)}(P(\alpha)).$$

It is clear that due to the re-sorting of consumers, friction-reducing policies change the demand, value and cost curves and these changes depend on the underlying micro-foundations. Importantly, the original demand, value and cost curves, considered in Einav et al. (2010) and Spinnewijn (2017), do not provide sufficient information for analyzing the market and welfare impact of such policies. However, the simple formulas in the Propositions (exploiting marginal policy changes) clearly indicate the key statistics underlying the overall effects we should anticipate:

Corollary 1 *In a competitive market with under-insurance, the marginal welfare gain from reducing information frictions is lower (and potentially negative) if (i) the mean friction value (i.e., $E_{P(\alpha)}(f)$) is higher, (ii) the re-sorting on costs (i.e., $-cov_{P(\alpha)}(c, f)$) is stronger and (iii) the re-sorting on surplus (i.e., $-cov_{P(\alpha)}(s, f)$) is weaker.*

To go beyond the local evaluations and provide further insights on how the primitives of the model, and the means and variances of the demand primitives in particular, impact positive and

normative outcomes under different policies, we present a series of simulations in Appendix Section D.¹¹ The simulations confirm the key insights of this theoretical analysis: (i) reducing the mean impact of frictions on willingness-to-pay for insurance always reduces insurance coverage, but reducing the variance of frictions can increase the demand for insurance when frictions suppress the demand of the marginal buyers. The latter occurs when mean surplus is relatively high such that equilibrium coverage is high as well; (ii) reducing the variance of frictions causes incremental adverse selection and reduces coverage more when the variance in costs is relatively high. The welfare implications tend to be in line with the implications for market function in an adversely selected market: equilibrium surplus decreases when equilibrium coverage decreases and vice-versa. The exception then holds when (iii) the variance of surplus is high relative to the variance of costs. The reason is that the positive matching effect of reduced frictions outweighs the negative equilibrium consequences of any incremental selection on costs, in line with the trade-off highlighted in Proposition 2.¹²

2.3.2 Risk-adjustment Transfers

The impact of demand frictions on equilibrium and welfare indicates their relevance for the evaluation of policies that target supply side frictions. We explore the importance of this interaction for cost subsidies and risk-adjustment transfers in particular. These policies are key features of US health reform, e.g., in the state exchanges set up under the ACA, as well as many other efforts to mitigate adverse selection and expand insurance coverage.

Risk-adjustment transfers subsidize the cost of providing insurance for an insurer based on the underlying risk of the insured individual. In practice, risk adjustment is implemented as a policy that facilitates transfers based on the realized or expected cost of the insured pool for each insurer.¹³ Introducing risk-adjustment in our stylized model, the expected cost to the insurer of providing insurance to individual i becomes

$$\tilde{c}_i(\beta) = c_i - \beta \times [c_i - Ec]$$

with $\beta \in [0, 1]$ and $\beta = 1$ capturing full risk-adjustment.

An increase in β makes the expected cost of providing insurance less dependent on the individ-

¹¹The importance of the relative variances of the demand components matter can be easily seen from rewriting the conditional covariances (as used in the Proposition 2) in terms of conditional variances of the demand primitives:

$$cov_P(x, f) = \frac{1}{2} [var_P(y) - var_P(x) - var_P(f)] \text{ for } x = c, s \text{ and } y = s, c.$$

¹²The (unconditional) correlations between the different demand components matter as well. A positive correlation between two demand components increases the conditional covariance between these two components. A negative correlation with a third demand component further increases the conditional covariance between the first two components.

¹³Whether risk adjustment compensates plans based on realized versus expected cost is an important question for the efficiency of incentives to insurers that trade off selection incentives against the power of cost reduction incentives conditional on enrollment. Geruso and McGuire (2016) study the issue in detail and we abstract from this tradeoff in our model and empirical implementation.

ual's risk type, but does not affect the ordering of individuals directly.¹⁴ In an adversely selected market, the average cost among the infra-marginal individuals unambiguously decreases for a given price. Hence, risk-adjustment transfers unambiguously reduce the equilibrium price and increase equilibrium coverage. Moreover, the more adversely selected the market is, the larger the impact of risk-adjustment transfers on equilibrium coverage. This indicates a first key interaction with information frictions as they can reduce selection on costs. Risk-adjustment transfers will affect the equilibrium by more the less plan selection is affected by demand frictions.

Since risk-adjustment transfers preserve the ordering of individuals' willingness-to-pay, the policy affects welfare only through the change in equilibrium coverage. The impact on welfare thus depends on the surplus among the marginal buyers in line with Proposition 1. This indicates a second key interaction with information frictions as the demand and supply frictions jointly determine whether the market is under- or over-insured. In an adversely selected market where information frictions reduce under-insurance, the presence of these frictions not only reduces the effectiveness of risk-adjustment transfers in increasing coverage, but also reduces the welfare gain from that increase.¹⁵ The following Proposition summarizes the potential effects:

Proposition 3 *A risk-adjustment policy β changes equilibrium coverage in a competitive market by*

$$Q'(\beta) = \eta^c \times [E_{\geq P}(c) - Ec].$$

The corresponding impact on welfare equals

$$\mathcal{W}'(\beta) = E_{P(\beta)}(s) Q'(\beta).$$

The above analysis highlights the important interaction between demand and supply side policies. Information policies can increase the effectiveness of risk-adjustment transfers and increase their impact on welfare. By the same token, the negative consequences of information policies through the increased adverse selection could be directly addressed through risk-adjustment transfers or any other policy that mitigates the increase in the equilibrium price. We confirm this complementarity between information policies and risk-adjustment in the simulations in Appendix Section D. We demonstrate that friction-reducing policies become more tenable, and can switch from 'bad' to 'good' as risk-adjustment is more effective. In particular, as the mean and variance of surplus increase relative to the mean and variance of costs in the population, the threshold

¹⁴Graphically, risk-adjustment transfers will flatten the cost curves relevant to the insurer relative to the demand curve. This is a key difference with a uniform subsidy, entailing a vertical shift of the original cost curves. In both cases, the ordering of individuals is maintained. This contrast with risk-rating where high-risk individuals pay a higher insurance premium than low-risk individuals. Risk-rating reduces sorting based on costs, but induces re-sorting based on frictions, analogue to our analysis of information policies.

¹⁵We again note that our risk adjustment framework assumes that a regulatory budget exists to fund risk adjustment transfers, and our welfare analysis does not explicitly consider the budgetary cost of the risk-adjustment policy equal to $\beta \times [E_{\geq P}(c) - Ec] \times Q(\beta)$. Though we do not do so here, it is not difficult to extend the model to account for different costs of funding.

of risk-adjustment necessary to make friction-reducing policies have a positive welfare impact is decreasing.

3 Empirical Application

We now move to our empirical application, which illustrates how the micro-foundations related to frictions, surplus, and costs can be measured and used to study (i) policies that impact choice and information frictions and (ii) insurer risk-adjustment transfers. We estimate these key micro-foundations using detailed proprietary data from a large self-insured employer covering more than 35,000 U.S. employees and 105,000 lives overall. The data include detailed administrative data on enrollee health care claims, demographics and plan choices as well as survey data, linked to the administrative data at the individual level, on consumer information and beliefs. The linked survey data allows us to go beyond previous empirical studies and distinguish between choice determinants and preference factors that are typically unobserved to researchers. This in turn permits the positive and normative analysis of both demand-side and supply-side policies. Though our empirical analysis studies one specific environment and population of consumers, it highlights how to connect the theoretical model just presented to data, and how to use those data together with an empirical framework to conduct important policy analyses.¹⁶

The data and estimation of consumer choice parameters we use are the same as that used in Handel and Kolstad (2015b), which performs an in depth study of consumer frictions and their implications for choice modeling in health insurance markets. That paper describes the data, empirical model, identification, estimation and structural choice parameter results in significant detail. Please see Handel and Kolstad (2015b) for a full treatment of that material. Here, we include a condensed summary of that content in Appendix F.

Key Micro-Foundations. The structural estimates from Handel and Kolstad (2015b) provide all the information we need to implement the approach developed in Section 2. We use the estimates to construct the micro-foundations that are key for determining market equilibrium and the impact of potential policy interventions.

Consumers in the empirical environment we study choose between two plan options, denoted j . The first option is a generous PPO option with zero cost-sharing, i.e. maximum risk protection. The second option is a high-deductible health plan (HDHP) with a \$3,750 family deductible and \$6,250 family out-of-pocket maximum that allows access to the same doctors in-network as the generous PPO option. The HDHP has an in-sample actuarial value of 78%, implying that, of all population expenses, consumers pay 22% of them. The HDHP plan also provides access to a health savings account (HSA) that provides some additional value to consumers by allowing them to make tax-free contributions to that plan that can be used to pay for health spending tax-free at any point (and accrue tax-free interest over time similar to a 401(k)).

¹⁶One directly relevant counterfactual market is a private insurance exchange offered by this large employer.

For each family k in the data we compute the perceived utility of choosing plan j :¹⁷

$$\begin{aligned}\widehat{U}_{kj} &= \int_0^\infty \widehat{f}_{kj}(z) \frac{-1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)\widehat{x}_{kj}(z)} dz \\ \widehat{x}_{kj}(z) &= W_k - P_{kj} - z + \mathbf{Z}'_k \widehat{\beta} \mathbf{1}_{j_i=HDHP} + \epsilon_{kj}\end{aligned}$$

Here, U_{kj} denotes consumers' constant absolute risk aversion (CARA) utility. X_k^A denotes observed heterogeneity (e.g. in age and income) for each family k . γ denotes the family-specific CARA risk aversion coefficient, which is estimated as a random coefficient with a normal distribution whose mean depends on X_k^A . f_{kj} denotes the ex ante rational expectations distribution of family out-of-pocket spending for family k and plan j , estimated using claims data. x_{kj} reflects a family's monetary equivalent value for each possible out-of-pocket health spending state realization (z). x depends on ex ante family wealth W , the premium paid P , and the amount of out-of-pocket health spending for one realization of uncertainty z . Additionally, it depends on $\mathbf{Z}'_k \widehat{\beta}$ which denotes family k 's additional willingness-to-pay for the *HDHP* relative to the *PPO* due to a collection of information frictions and perceived hassle costs that are measured with the linked survey in Handel and Kolstad (2015b). β is an estimated vector of coefficients that tells us how much each possible friction in the vector Z_k impacts consumer willingness-to-pay. For most frictions measured Z is a binary indicator of whether the consumer has limited information on a given dimension, though in certain cases Z is a real number reflecting the extent of a certain friction (e.g. the number of additional hours of hassle costs one incurs when enrolling in the HDHP, relative to the PPO). ϵ_{kj} is a family-specific idiosyncratic preference for each plan j .

We map these estimated utilities into our theoretical framework and define the willingness-to-pay for the PPO, relative to the HDHP, as the difference in certainty equivalents implied by the above utility model:

$$w_k = CE_{k,PPO} - CE_{k,HDHP} \quad (2)$$

Here, $CE_{k,j}$ is the certain financial payment that gives family k utility U_{kj} , equivalent to choosing plan j given the present frictions. The relative willingness-to-pay w_k is the empirical analog to w in Section 2. Figure 3 presents its distribution in the observed environment. This distribution determines the demand curve in our upcoming analysis and is plotted for families (employees covering 2+ dependents), who comprise the majority of our primary sample. Consumer willingness to pay for the PPO is high, but there is substantial heterogeneity in willingness to pay across families. To assess the main drivers of the observed heterogeneity, we decompose the willingness-to-pay into the different demand primitives, following the approach in Section 2.

¹⁷Note that we assume that consumers make active choices and have no default option (i.e., no inertia) to focus our analysis, though Handel and Kolstad (2015b) includes estimates of consumer inertia. Our analysis focuses on information frictions, but can be naturally extended to assess the impact of reducing inertia, in isolation or joint with the reduction of information frictions.

[FIGURE 3 ABOUT HERE]

First, the coefficient estimates on each friction allow us to assess the combined impact of all frictions on the willingness to pay for each family. To construct the empirical analog of the friction value f from Section 2, we simply use:

$$f_k = -\mathbf{Z}'_k \widehat{\beta}_{j_t=HDHP} \quad (3)$$

The obtained value describes how much the frictions present shift willingness-to-pay relative to an equivalent frictionless consumer. Figure 3 presents the smoothed distribution of the combined impact of all frictions on willingness to pay for less generous coverage relative to more generous coverage (i.e., $-f$). As the figure illustrates, the information frictions have a high mean impact of shifting consumers towards more generous coverage (\$1787, see Table F3) as well as substantial heterogeneity (standard deviation of \$1304). Thus, our empirical environment corresponds most closely to the case with high mean friction impact and high friction heterogeneity discussed in Section 2.

Second, from the cost model (described in detail in Appendix B), we obtain an estimate of the distribution of total expenses for each family. Appendix Figure F2 plots the distribution of expected total expenses for each family: as is typical this is a fat-tailed distribution similar to a lognormal distribution with a fairly large degree of consumer heterogeneity and a high level of mean spending. Using the plan characteristics of the offered PPO and HDHP plans, we map the distribution of expenses for each family into expected insurer costs from providing each plan j to family k . Define $c_{k,PPO}$ as the expected insurer costs for the PPO and $c_{k,HDHP}$ as insurer costs for just the HDHP (i.e. the baseline plan). The difference between the two equals the supplemental insurer cost, which is the empirical analog to c in Section 2:

$$c_k = c_{k,PPO} - c_{k,HDHP}.$$

Figure 3 plots the smoothed distribution of the expected insurer costs from providing the supplemental coverage, c , for families in our primary sample. The figure reveals substantial heterogeneity in insurer costs. The consumer's out-of-pocket maximum of the HDHP, however, imposes an upper bound on the supplemental insurer costs, showing up as a spike in the distribution.

Finally, having determined willingness-to-pay, friction impact and insurer costs, we can compute incremental welfare from additional risk protection (the empirical analog to surplus s) as the difference between 'true' insurance value $v_k = w_k - f_k$ and actual relative cost c_k :

$$s_k = v_k - c_k$$

Figure 3 also presents the distribution of surplus from risk protection for the PPO relative to the HDHP. The distribution of surplus is skewed towards 0, since many consumers are estimated to be near risk-neutral, though there is a non-trivial group of consumers with substantial positive surplus. Overall, the mean and variance of this surplus are substantially lower than the means and

variances of the cost distribution and the friction distribution.

In the context of our theoretical analysis, our empirical environment is one with high mean and variance of frictions, low mean and variance of surplus, and medium to high mean and variance in expected yearly costs. As a result, as we saw in that section, we expect that friction-reducing policies will lead to substantial unraveling in the absence of complementary risk-adjustment.

[TABLE 1 ABOUT HERE]

Table 1 presents the correlations between these micro-foundations for families in our primary sample. The first thing to note is that the impact of frictions is relatively uncorrelated with surplus from risk protection, cost, and true value for more generous coverage. It is highly correlated with willingness-to-pay, since frictions are large in magnitude and feed directly into willingness-to-pay. Surplus from risk protection is highly correlated with cost and with true plan value, but less correlated with willingness-to-pay due to the presence of frictions. Cost is almost perfectly correlated with true value, because of limited heterogeneity in risk aversion, while frictions are the strongest correlate of willingness-to-pay. Frictions are thus an important determinant of demand in our environment, as are costs, but costs become much more highly correlated with willingness-to-pay when frictions are removed.

Market Setup. The primary counterfactual market we consider is, as described in Section 2, a competitive market for supplemental insurance that moves consumers from universal baseline coverage (represented by the HDHP in our empirical environment) to more generous overall coverage (represented by the PPO). We assume that an individual mandate is enforced, such that individuals enroll in either the public baseline coverage, or that coverage plus the supplemental coverage (for this market, this is similar to saying the public coverage is provided for free).

We make the important assumption that the relative information frictions we estimate for our two empirical plans map directly to the relative information frictions that consumers have for supplemental coverage relative to the baseline coverage. This assumption would be violated, e.g., if competing insurers worked harder to either provide or obscure information relative to what the firm in our empirical environment does. This analysis should thus be viewed as a stylized analysis that highlights the potentially nuanced implications of friction-reducing policies together with risk-adjustment policies, rather than an analysis that makes specific predictions of what will happen in a particular regulated marketplace.¹⁸

We study a range of demand-side policies that reduce consumer choice frictions and supply-side policies that impact the costs insurers face for different consumers. Using our structural estimates of frictions, surplus and costs we construct (i) demand curve (ii) welfare-relevant value curve and (iii) average and marginal cost curves for each policy scenario.

¹⁸In Appendix E, we also present some results for the class of markets studied in Handel et al. (2015) where insurers compete to offer two types of insurance policies simultaneously and an individual mandate is in place requiring consumers to buy one of the two types of policies. Construction of demand and value for incremental coverage is the same as in the primary markets studied in the main text, but construction of average and marginal cost curves is different, reflecting the internalization of costs by the lower coverage plans in that setup.

The demand curve reflects consumer willingness-to-pay for more generous coverage in a given policy environment. This willingness-to-pay is the same regardless of whether it is a market for supplemental add-on coverage or a market where insurers offer both types of plans. Of the two policies we consider here — those that reduce information frictions and insurer risk adjustment transfers — only the former impacts consumer demand. As a result, counterfactual consumer willingness-to-pay for each plan j given a specific information friction reduction policy α is:

$$\begin{aligned}\widehat{U}_{kj}(\alpha) &= \int_0^\infty \widehat{f}_{kj}(z) \frac{-1}{\gamma_k(X_k^A)} e^{-\gamma_k(X_k^A)\widehat{x}_{kj}(\alpha,z)} dz \\ \widehat{x}_{kj}(\alpha, z) &= W_k - P_{kj} - z + (1 - \alpha)\mathbf{Z}'_k\widehat{\beta}\mathbf{1}_{j_t=HDHP} + \epsilon_{kj}\end{aligned}$$

Thus, when $\alpha = 0$ all information frictions are present and consumer demand is composed of estimated willingness-to-pay for each plan in our given environment. When $\alpha > 0$ then information frictions are reduced by some fraction, up to the case when $\alpha = 1$ and 100% of frictions are removed. In our upcoming analysis, we investigate a space of policies corresponding to values of α between 0 and 1. The level of α can be thought of as a reduced form representation of different policy combinations that reduce consumer choice frictions (e.g., information provision, decision support, or smart defaults).¹⁹

Willingness-to-pay for the PPO, relative to the HDHP, for family k given the friction-reducing implications of α equals:

$$\tilde{w}_k(\alpha) = C\widehat{E}_{k,PPO}(\alpha) - C\widehat{E}_{k,HDHP}(\alpha). \quad (4)$$

This simplifies to $\tilde{w}_k(\alpha) = w_k - \alpha \times f_k$ as in Section 2. The corresponding relative demand curve for the *PPO* relative to the *HDHP* equals:

$$D(P; \alpha) = Pr(\tilde{w}_k(\alpha) \geq P)$$

Here, P is the price of supplemental coverage that moves the consumer from the baseline HDHP plan to combined coverage represented by the PPO plan.

The welfare-relevant value curve $V(P; \alpha)$ reflects the value of additional coverage in an environment with no information frictions v_k (i.e., $v_k = \tilde{w}_k(1)$), conditional on *the same ordering* of consumers as $D(P; \alpha)$:

$$V(P; \alpha) = E[v|\tilde{w}_k(\alpha) = P]$$

The empirical value curve only coincides with the demand curve when $\alpha=1$: for other values of α each consumer's true value is the same, but the ordering of consumers along the value curve is different, since the demand curve reflects both value and information frictions.²⁰

¹⁹Though we do not quantify the empirical impact of actual friction-reducing policies in this paper, one could in principle study values of α linked to specific empirical measures and/or policy changes.

²⁰As mentioned before, the construction of $V(P; \alpha)$ embeds the assumption that the estimated demand impacts of

The average and marginal cost curves relevant to the insurer are determined by the insurer costs and the insurer risk-adjustment transfers, but also depend on the underlying preferences and information frictions (due to the sorting effect). Risk-adjustment transfers compensate insurers for a share β of the difference in costs for the selection of families buying insurance and the average cost in the population. In the market for supplemental insurance the marginal cost curve is defined as follows for a given policy combination (α, β) :

$$MC(P; \alpha, \beta) = E[c_k | \tilde{w}_k(\alpha) = P] - \beta E[c_k | \tilde{w}_k(\alpha) = P] - (AC_{pop, PPO} - AC_{pop, HDHP}),$$

where $\beta = 1$ denotes perfect risk-adjustment. This is the insurer MC curve given risk-adjustment: the true marginal cost curve, *which is the cost curve relevant for welfare analysis*, is defined as the insurer marginal cost curve where $\beta = 0$ (i.e., $MC(P; \alpha, 0)$ for each P). The average cost curve $AC(P; \alpha, \beta)$ simply traces out the average of supplemental costs for those with willingness to pay greater than or equal to P :

$$AC(P; \alpha, \beta) = E[c_k | \tilde{w}_k(\alpha) \geq P] - \beta E[c_k | \tilde{w}_k(\alpha) \geq P] - (AC_{pop, PPO} - AC_{pop, HDHP})$$

The insurer cost curves depend on α because, as frictions are reduced, costs become a more prominent driver of demand. Consequently, the correlation between costs and willingness to pay becomes higher, leading to different costs curves as a function of quantity demanded at a given relative price. The insurer cost curves also depend on β , the insurer risk-adjustment transfers, because as risk-adjustment transfers are implemented between insurers the contribution of a given consumer to plan cost is mitigated by transfers and the curves become flatter. Equilibrium in the market occurs at the lowest value of P such that $P = AC(P; \alpha, \beta)$, under a set of regularity conditions which we assume hold here.²¹

Once we have determined the equilibrium outcome in each market, we compute incremental consumer welfare from more generous coverage as:

$$\sum_k s_k 1[\tilde{w}_k(\alpha) \geq P]$$

For a given equilibrium allocation and price P , the welfare loss relative to the first-best, where

observed information frictions are not welfare-relevant once a consumer is actually allocated to a given plan. For some of the frictions we study (e.g., information about provider networks) this assumption seems very reasonable, while for others (e.g., perceived hassle costs) this is less clear. It is straightforward to alter the definition of V for different underlying models mapping revealed willingness-to-pay and measured frictions to welfare-relevant valuations. See Handel and Kolstad (2015b) for an in depth discussion of the welfare implications for each specific friction studied.

²¹We also note here that, because there is only one type of non-horizontally differentiated priced plan, risk-adjustment implies a transfer into (or out of) this supplemental market if the market is adversely (advantageously) selected. This is a feasible policy approach both in theory and practice (see the the discussion in e.g., Handel et al. (2015) or Mahoney and Weyl (2017) for greater detail). Finally, we note that for the alternative market setup where both coverage tiers are competitively offered, construction of the average and marginal costs curves is different than for the supplemental market described here. See Appendix E for a lengthier discussion.

everyone enrolls in more comprehensive coverage (i.e., $s > 0$), is:

$$\sum_k s_k 1[\tilde{w}_k(\alpha) < P]$$

Using this metric, in the next section we compare the welfare impact of different friction-reducing and risk-adjustment policies, both relative to other candidate policies and relative to a first-best.

Empirical Results. In our empirical application, we first evaluate the positive and normative implications of friction-reducing policies on their own and then discuss the impact of these policies conditional on different levels of risk-adjustment effectiveness. We focus on the Einav et al. (2010) style market for supplemental coverage, which provides incremental coverage relative to the HDHP baseline plan.²² We present results only for the family coverage tier, who comprise the majority of our sample and form a natural population for a community rated market (since typically firms can vary premiums w/ number of enrollees).²³

Information frictions impact both the number of individuals buying each type of plan and the sorting of individuals across plans. Therefore, we expect both the level and slope of the demand, cost, and value curves to change as α changes. Figure 4 presents these sets of curves graphically for full ($\alpha = 0$), half ($\alpha = .5$) and no ($\alpha = 1$) choice frictions. Recall that when $\beta = 0$ and there is no risk-adjustment, as in these figures, the true marginal cost curve for consumers is the same as the insurer marginal cost curve. Note also that the value and marginal cost curves correspond to the same ordering of individuals as the demand curve for each scenario.

The leftmost panel in Figure 4, which replicates the demand, value, and cost curves as estimated in our environment (with all frictions present), illustrates some key implications of our estimates. First, the frictions present in our environment drive a substantial wedge between the demand curve and welfare-relevant value curve: the demand curve lies well above the value curve, indicating that consumers on average over-value the more comprehensive *PPO* plan relative to the *HDHP*. This is true along the entire demand curve, even for consumers with a relatively low willingness to pay for the supplemental coverage. Second, it is clear from the charts that the surplus of the supplemental coverage is quite small, especially relative to the impact of frictions on willingness-to-pay. In each figure, surplus is represented by the wedge between the marginal cost curve and the welfare-relevant value curve and corresponds to the risk-premia consumers are willing to pay to be in the *PPO* as opposed to the *HDHP*. While the average cost curve is downward sloping — a necessary condition for adverse selection — the slope is relatively flat. This indicates that, when full frictions are present, marginal enrollee costs to the *PPO* are not substantially different than those of infra-marginal enrollees and there is limited scope for adverse selection.

[FIGURE 4 ABOUT HERE]

²²We present the empirical results for exchange-style markets with two priced plans in Appendix E.

²³For all results, we present a version of our estimates that fits the non-parametric curves with splines: upon request we have completed and can provide a linearized version (as in Einav et al. (2010)), which is more restrictive, and a fully non-parametric version, which is less restrictive.

Table 2 presents the positive market equilibrium results associated with different policy combinations. The first column, for $\beta = 0$ gives the results for the cases of different friction-reducing policies when there is no insurer risk-adjustment (as shown in Figure 4). In all cases, since the value curve lies about the consumer marginal cost curve, 100% of consumers should be allocated to the *PPO* from a social perspective. In our conclusions, we return to these results and discuss un-modeled factors that would change this first-best allocation, such as moral hazard.

For the case of full frictions ($\alpha = 0$) the predicted market equilibrium outcome (the one crossing point between the demand and average cost curves) is 84.6% enrolled in the *PPO* and 15.4% enrolled in the *HDHP*. The price paid for supplemental coverage in equilibrium equals $P = \$5,551$. For the case of half frictions ($\alpha = 0.5$, Figure 4), 73.4% buy the *PPO* and 26.6% buy the *HDHP* in equilibrium, with a relative premium difference of $P = \$5,741$. When the impact of frictions are reduced by 50% there is only limited incremental adverse selection against the *PPO*, with market share declining and the relative premium rising.

When all frictions are removed ($\alpha = 1$, Figure 4) the demand curve and value curve are equivalent, with demand shifting downward relative to the case where frictions are present. In addition, the marginal and average cost curves become steeper reflecting the sorting effect as consumer marginal costs are much more highly correlated with consumer demand. The market equilibrium reflects an almost complete unraveling of the market due to adverse selection: 9.1% of consumers buy the *PPO*, 90.9% buy the *HDHP* and the relative premium is $P = \$6,250$.

Both the level and sorting effects lead to the unraveling of the market as information frictions are reduced in our environment. The level effect can be seen clearly in Figure 4 above, as the demand shifts down substantially as frictions are reduced (also for the marginal consumers). The sorting effect can be seen clearly in Figure 5: as frictions are reduced the average cost curve becomes steeper, implying that the correlation between consumer costs and demand is increasing. Table 1 shows that this correlation increases from 0.508 to 0.999 as frictions are reduced to non-existent. In essence, the presence of information frictions drives a gap between demand and welfare-relevant valuation, and the correlation of those frictions with costs determines if removing frictions has a marked sorting effect. In our case, frictions are not particularly highly correlated with costs, so when they are present they have a substantial impact on the ordering of willingness-to-pay for more insurance.

The bottom portion of Table 2 presents the welfare implications of friction-reducing policies. In our environment, where consumers benefit from more risk protection (assuming no corresponding efficiency loss from increased moral hazard), welfare is generally decreasing as the market unravels and enrollment in the more generous *PPO* plan goes down (this is not necessarily true because of the improved matching as discussed before). Our welfare results show that, relative to the status quo environment, when frictions are reduced by 50% consumers are worse off by an average of \$16.04 (35% of mean total surplus) per person. When frictions are fully removed and the market unravels, consumers are on average \$47.01 (99% of mean total surplus) worse off per person. This is a meaningful drop in welfare for a policy is typically thought to benefit consumers.

[FIGURE 5 ABOUT HERE]

[TABLE 2 ABOUT HERE]

One way to counter these welfare losses are risk-adjustment transfer policies. We demonstrate the impact of risk-adjustment policies spanning $\beta = 0$ to $\beta = 1$ conditional on $\alpha = 1$, or when frictions are already fully removed. Figure 6 presents the demand curve for $\alpha = 1$ (equivalent to the value curve) and three average cost curves, corresponding to the cases of $\beta = 0$, $\beta = 0.5$, and $\beta = 1$. From the figure, it is clear that as risk-adjustment becomes stronger, the average cost curve becomes flatter, becoming completely flat when $\beta = 1$ and all consumers have the same cost from the insurer's perspective. It is also apparent that as risk-adjustment becomes more effective, the market share of the *PPO* plans increase, and the market equilibrium moves towards the first-best of 100% *PPO* enrollment. Table 2 presents the resulting market shares and premiums: for the cases of $\beta = 0$, $\beta = 0.5$, and $\beta = 1$ the resulting market shares when $\alpha = 0$ are 9.1%, 51.6%, and 63.5% respectively. The relative premiums between the two tiers of plans are \$6,250, \$5,964, and \$5,315 respectively. Thus, conditional on frictions being fully removed, risk-adjustment has a substantial impact of reducing premiums in the *PPO* relative to the *HDHP*, and increasing market share in the *PPO*. Welfare in the market is also increasing as insurer risk-adjustment policies become more effective. When frictions are fully removed, risk-adjustment that is 50% effective increases welfare by 19% of mean total surplus (\$8.71) per person on average. When risk-adjustment is 100% effective, welfare increases by 39% of mean total surplus (\$17.67) per person on average.

[FIGURE 6 ABOUT HERE]

Figure 6 also presents the same curves for these three risk-adjustment policies, for the case of $\alpha = 0$ (our observed environment). Here, though the directional impacts of stronger risk-adjustment on plan market shares and relative premiums are the same as when $\alpha = 1$, the incremental effect is much weaker because the frictions present in the environment already reduce adverse selection to a large extent. The quantity in the *PPO* increases from 84.2% to 88.5% as β goes from 0 to 1, with the relative price decreasing from 5,551 to 5,315. The corresponding impact on welfare is again positive, but small. Welfare increases by 9% of mean total surplus (\$4.30) per person on average.

These findings make clear that the marginal impact of either (i) friction-reducing policies or (ii) insurer risk-adjustment transfers depends crucially on the effectiveness of the other policy within any given environment. One important implication of this is that policymakers considering policies to improve consumer decisions may want to simultaneously strengthen insurer risk-adjustment policies in order to prevent incremental adverse selection. This is especially true in cases like our empirical environment, where the mean and variance of surplus are low relative to the mean and variance of costs.

Figure 7 plots market equilibrium quantities, prices, and welfare outcomes for all combinations of policies $\alpha \in [0, 1] \times \beta \in [0, 1]$. Select numbers from the three panels in the figure are reported in Table 8. The key insight across all three panels in the Figure is that effective risk-adjustment becomes

increasingly impactful and important as information frictions are reduced. For low to medium values of α , where substantial choice frictions are still present, more effective risk-adjustment has only a minimal impact on market outcomes and welfare. This is because the average cost curve is already quite flat for low values of α , so there is not much scope for risk-adjustment to further change market outcomes by resorting consumers and further flattening the cost curve. However, for high values of α , where the cost curves are steeper and preferences have been shifted towards the *HDHP* via the level effect, risk-adjustment has an immediate and strong effect by flattening the cost curve, reducing adverse selection and improving market outcomes. Simply put, if consumer choices are less responsive to a consumer’s specific cost, decoupling insurer pricing from individual specific risk has less of an impact.

[FIGURE 7 ABOUT HERE]

The rightmost panel in Figure 7 and the bottom panel of Table 8 show the welfare impact of possible policy combinations in the $\alpha - \beta$ space. Risk-adjustment policies have a large incremental impact when friction-reducing policies are very effective: when $\alpha = 1$ moving β from 0 to 1 improves welfare by \$17.67 per person on average, while when $\alpha = 0$ the same movement in β improves welfare by \$4.30 per person on average. For $\alpha = 0.2$, $\beta = 1$ still leads to a welfare improvement relative to the status quo, while for values $\alpha = 0.5$ and above no degree of risk-adjustment improves welfare relative to the baseline case.

As a final note, we emphasize that this empirical analysis reflects the case where there is low mean consumer surplus from incremental insurance and low surplus variance, relative to both the degree of frictions in the market and the variance in projected costs. As a result, as frictions are removed, the market unravels relatively quickly because costs feed back into premiums but lower cost consumers don’t have high enough true surplus to justify the purchase of incremental insurance when frictions are reduced. In different insurance environments, the mean and variance of surplus may be larger (e.g., if there is no out-of-pocket maximum or consumers are more risk averse than those here) which, as our simulations in Section 2 reveal, may lead friction reducing policies to have positive impacts on their own. In such cases, friction-reducing policies can and should be implemented even if effective risk-adjustment is not available.

4 Conclusion

In this paper we set up a general framework to study insurance market equilibrium and the welfare that results for environments where limited information distorts consumer plan choices. Understanding the relationship between the key micro-foundations – (i) surplus from risk protection (ii) the impact of frictions on willingness-to-pay and (iii) consumer/insurer costs – is essential for making policy decisions. We use this framework to investigate demand-side policies that reduce consumer information frictions, thereby helping consumers make better plan choices, and insurer risk-adjustment transfers, a supply-side policy designed to mitigate adverse selection by dampening the relationship between consumer costs and insurer costs.

Our theoretical framework and empirical application highlight the subtleties that determine when policies to reduce consumer frictions will be welfare increasing or welfare decreasing. Crucially, the impact of these policies depends not only on the distributions of micro-foundations in a market, but also on how effective complementary supply-side policies, such as insurer risk-adjustment transfers are. If insurer risk-adjustment policies are not shown to be highly effective (see e.g., Brown et al. (2014)), then policymakers may want to be more conservative in implementing policies that heavily reduce the impact of information frictions in the market. This is especially true in cases where the mean and variance of costs are high relative to those of consumer surplus. However, when considering more horizontally differentiated markets with strong variation in consumer surplus, the opposite could be true. These insights are important for policymakers thinking about implementing policies such as information provision, plan recommendations, and smart defaults, all of which are being currently considered by different insurance market regulators.

Our empirical example illustrates how our theoretical framework can be implemented empirically in different contexts with distinct micro-foundations. Previous work suggests that these micro-foundations could be meaningfully different across insurance market environments. For example, a range of papers show meaningful consumer choice frictions in Medicare Part D (see for example Abaluck and Gruber (2011) or Ketcham et al. (2012)), where the mean and variance of costs is lower than in our market (because it insurers only prescription drugs) and risk-adjustment may be very effective (because of the predictability of drug use). In that case, our framework suggests that friction-reducing policies are more likely to be welfare improving than in the empirical environment we investigate in this paper. Of course, the relevant micro-foundations must be measured in each context to directly apply our framework, though our results demonstrate methods to do so as well as the feasibility. For parsimony our discussion focused on the case where frictions push consumers towards more generous coverage, which has been found in several studies of choice in employer-sponsored insurance settings (e.g. Handel and Kolstad (2015b), Handel (2013) and Bhargava et al. (2017)). Our framework can also be applied to the reverse case where frictions push consumers towards purchasing less coverage, which research shows may be relevant in certain contexts such as the subsidized ACA exchanges (e.g. Finkelstein et al. (2017)).

Our framework contains a range of stylized assumptions that could impact the conclusions in any given context. We assume perfect competition: as Mahoney and Weyl (2017) show, imperfect competition can have subtle implications for policy recommendations in selection markets. Additionally, our approach maintains quite stylized assumptions about the potentially endogenous relationship between the extent of competition in the market and consumer information. It is possible that the extent of limited information in any given setting is partially related to the degree of competition and/or the extent of risk-adjustment policies, an area that we believe is an interesting topic for future work. We also abstracted away from consumer moral hazard, to clearly focus on the other micro-foundations in the market. Though the relationships we explore would generally be robust to including moral hazard in the model, the mean and variance of that price sensitivity could have important implications for whether increasing coverage is a desirable social goal.

References

- Abaluck, J. and J. Gruber**, “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program,” *American Economic Review*, 2011, *101* (4), 1180–1210.
- Akerlof, G.**, “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 1970, *84* (3), 488–500.
- Alcott, H. and D. Taubinsky**, “Evaluating Behaviorally-Motivated Policy: Evidence from the Lightbulb Market,” *The American Economic Review*, 2015, *105* (8), 2501–2538.
- Azevedo, E. and D. Gottlieb**, “Perfect Competition in Markets with Adverse Selection,” *Econometrica*, 2017, *85* (1), 67–105.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein**, “Behavioral Hazard in Health Insurance,” *Quarterly Journal of Economics*, 2015, *130*, 1623–1667.
- Bernheim, D., A. Fradkin, and I. Popov**, “The Welfare Economics of Default Options in 401(k) Plans,” *The American Economic Review*, 2015, *105* (9), 2798–2837.
- 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.
- Bhargava, S., G. Loewenstein, and J. Sydnor**, “Choose to Lose: Health Plan Choices from a Menu with Dominated Options,” *Quarterly Journal of Economics*, 2017, *132* (3), 1319–1372.
- Bronnenberg, Bart J., Jean-Pierre Dube, Matthew Gentzkow, and Jesse M. Shapiro**, “Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium,” August 2014. University of Chicago Working Paper.
- Brot-Goldberg, Z., A. Chandra, B. Handel, and J. Kolstad**, “What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics,” *Quarterly Journal of Economics*, 2017, *132* (3), 1261–1318.
- Brown, J., M. Duggan, I. Kuziemko, and W. Woolston**, “How Does Risk Selection Respond to Risk Adjustment? Evidence From The Medicare Advantage Program,” *American Economic Review*, 2014, *104* (10), 3335–3364.
- Bundorf, Kate, Jonathan Levin, and Neale Mahoney**, “Pricing and Welfare in Health Plan Choice,” *American Economic Review*, 2012, *102* (7), 3214–3248.
- Chetty, R. and A. Finkelstein**, “Social Insurance: Connecting Theory to Data,” 2013. in *Handbook of Public Economics*, vol 5, 111 - 193.
- Cutler, D. and S. Reber**, “Paying for Health Insurance: The Tradeoff Between Competition and Adverse Selection,” *Quarterly Journal of Economics*, 1998, *113* (2), 433–466.
- Dixit, A. and V. Norman**, “Advertising and Welfare,” *The Bell Journal of Economics*, 1978, *9* (1), 1–17.
- Einav, L., A. Finkelstein, and M. Cullen**, “Estimating Welfare in Insurance Markets Using Variation in Prices,” *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.
- Finkelstein, A., N. Hendren, and M. Shepard**, “Subsidizing Insurance for Low-Income Adults: Evidence from Massachusetts,” August 2017. NBER Working Paper No. 23668.

- Geruso, M. and T. McGuire**, “Tradeoffs in the Design of Health Plan Payment Systems: Fit, Power, and Balance,” *Journal of Health Economics*, 2016, 47, 1–19.
- Hackmann, M., J. Kolstad, and A. Kowalski**, “Adverse Selection and an Individual Mandate: When Theory Meets Practice,” *American Economic Review*, 2015, 105 (3), 1030–1060.
- Handel, B.**, “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 2013, 103 (7), 2643–2682.
- **and J. Kolstad**, “Getting the Most from Marketplaces: Smart Policies on Health Insurance Choice,” *Hamilton Project Discussion Paper*, 2015, (8).
- **and –**, “Health Insurance For Humans: Information Frictions, Plan Choice, and Consumer Welfare,” *American Economic Review*, 2015, 105 (8), 2449–2500.
- **, I. Hendel, and M. Whinston**, “Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk,” *Econometrica*, 2015, 83 (4), 1261–1313.
- 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.
- Mahoney, N. and G. Weyl**, “Imperfect Competition in Selection Markets,” *Review of Economics and Statistics*, 2017, 99 (4), 637–651.
- Polyakova, M.**, “Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D,” *American Economic Journal: Applied Economics*, 2016, 8 (3), 165–195.
- Rees-Jones, A. and D. Taubinsky**, “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment,” *The Review of Economic Studies*, forthcoming.
- 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.
- Spinnewijn, J.**, “Heterogeneity, Demand for Insurance and Adverse Selection,” *American Economic Journal: Economic Policy*, 2017, 9 (1), 308–343.
- Train, K.**, *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge University Press, 2009.
- Veiga, A. and G. Weyl**, “Product Design in Selection Markets,” *Quarterly Journal of Economics*, 2016, 131 (2), 1007–1056.
- Weyl, G. and A. Veiga**, “Pricing Institutions and the Welfare Cost of Adverse Selection,” *American Economic Journal: Microeconomics*, 2017, 9 (2), 139–148.

Tables and Figures

Correlations					
Key Micro-Foundations					
Primary Estimates					
	Friction f	Surplus s	Cost c	WTP	Value
Friction f	1	-0.0082	0.0249	0.8485	0.0269
Surplus s	-0.0082	1	0.8418	0.4418	0.8211
Cost c	0.0249	0.8418	1	0.5075	0.9993
WTP ($f + s + c$)	0.8485	0.4418	0.5075	1	0.5062
Value ($s + c$)	0.0269	0.8211	0.9993	0.5062	1

Table 1: This table presents key correlations between (i) impact of frictions on *PPO* willingness to pay (ii) incremental surplus from *PPO* risk protection (iii) expected marginal *PPO* health spending for insurer (iv) willingness to pay for *PPO* and (v) true relative *PPO* value. Results are presented for families (covering at least a spouse and dependent) who comprise over 50% of our primary sample.

Positive Policy Impacts

	$\beta = 0$	$\beta = .2$	$\beta = .5$	$\beta = .8$	$\beta = 1$
Quantity <i>PPO</i>					
$\alpha = 0$	84.6%	85.5%	87.1%	88.0%	88.5%
$\alpha = .2$	81.9%	83.3%	84.6%	85.7%	86.4%
$\alpha = .5$	73.7%	76.1%	78.6%	80.9%	82.0%
$\alpha = .8$	51.7%	59.4%	68.0%	72.0%	74.1%
$\alpha = 1$	9.1%	34.7%	51.6%	59.0%	63.5%

Price of Supplemental Coverage *P*

$\alpha = 0$	\$5,551	\$5,498	\$5,425	\$5,358	\$5,315
$\alpha = .2$	\$5,611	\$5,544	\$5,452	\$5,368	\$5,315
$\alpha = .5$	\$5,741	\$5,643	\$5,507	\$5,385	\$5,315
$\alpha = .8$	\$6,035	\$5,835	\$5,596	\$5,418	\$5,315
$\alpha = 1$	\$6,250	\$6,014	\$5,694	\$5,452	\$5,315

Welfare Impact*

	$\beta = 0$	$\beta = .2$	$\beta = .5$	$\beta = .8$	$\beta = 1$
$\alpha = 0$	0	0.97	2.34	3.30	4.30
$\alpha = 0.2$	-4.96	-3.44	-1.98	-0.58	0.40
$\alpha = 0.5$	-16.04	-13.92	-11.15	-7.97	-6.36
$\alpha = 0.8$	-36.63	-31.29	-24.21	-20.23	-17.95
$\alpha = 1$	-47.01	-45.13	-38.30	-33.15	-29.34

*Relative to ($\alpha = 0, \beta = 0$)

Table 2: The first two sections of this table present the market outcomes in prices and quantities for different policy combinations of (i) friction-reducing policies and (ii) insurer risk-adjustment transfers. The third panel presents the relative welfare impact of different policies; policies are compared to information frictions and zero risk adjustment ($\alpha = 0$ and $\beta = 0$).

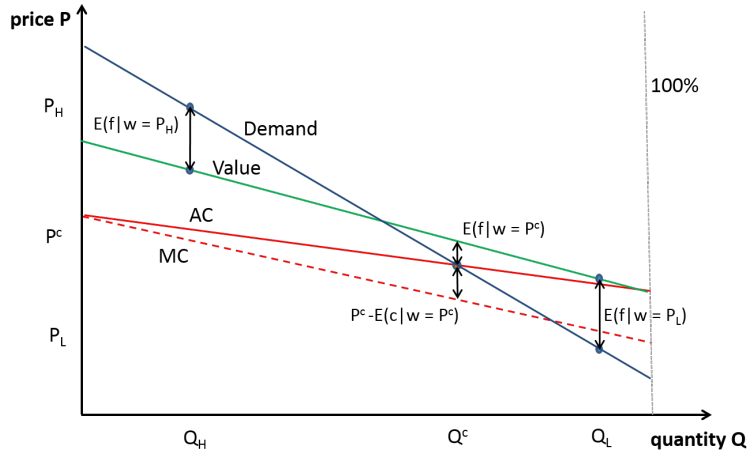


Figure 1: Demand, value and cost curves in an adversely selected market with heterogeneous frictions

Notes: The figure shows the share of individuals buying insurance $Q = D(P)$ on the horizontal axis, for each price P on the vertical axis. The figure also shows the expected costs for the marginal and infra-marginal buyers and the expected value for the marginal buyers at that price P , again on the vertical axis. Information frictions drive a wedge between the demand curve and the value curve in two ways. First, for a uniform friction $f_i = \bar{f}$, the value curve is parallel to the demand curve, $E_P(v) = P - \bar{f}$. Second, heterogeneous demand frictions $f_i = \bar{f} + \varepsilon_i$ cause the value curve to be a counter-clockwise rotation of the demand curve when the friction variation is independent. Individuals with higher willingness-to-pay tend to overestimate the value of insurance more while individuals with sufficiently low willingness-to-pay underestimate the insurance value despite a positive average friction. This causes the average friction value $E_P(f)$ to become negative for consumers with low willingness-to-pay in the Figure. The vertical difference between the value curve and the marginal cost curve for a given level of market coverage equals the expected surplus for the marginal buyers. The Figure plots the case where value always exceeds cost. Total welfare corresponds to the difference between the value curve and the marginal cost curve for all individuals buying insurance.




Figure 2: Sorting effect of friction-reducing policies: value and frictions among the marginal consumers

Notes: The figure shows the combinations of true values v and friction values f for which an individual buys insurance. A downward sloping curve implied by $v + (1 - \alpha) f = P$ separates the group of insured and uninsured. This curve flattens due to an information policy; the individuals who start buying insurance have higher true value than the individuals who stop buying insurance. The information policy thus necessarily increases the expected true value $E_{>P}(v)$ for a given share of buyers. Or equivalently, the covariance between true and friction value among the marginal buyers, $cov_P(v, f)$, is necessarily negative.




Figure 3: This figure presents the smoothed estimated distributions of key consumers micro-foundations in our empirical application. Estimates are presented for families (employees covering 2+ dependents), who comprise the majority of our sample and who are the focus of our upcoming counterfactual market analysis. The figure presents the distributions of (i) consumer willingness-to-pay for the PPO relative to the HDHP (top left) (ii) total impact of frictions on willingness-to-pay for the HDHP relative to the PPO (top right) (iii) expected supplemental insurer costs from the PPO relative to the HDHP (bottom left) and (iv) surplus from risk protection for the PPO relative to the HDHP (bottom right).




Figure 4: From left to right, these figures show (i) market equilibrium including information frictions (ii) market equilibrium with partial information frictions ($\alpha = 0.5$) and (iii) market equilibrium without information frictions.




Figure 5: Average Cost Curves with Varying Levels of Information Frictions.




Figure 6: From left to right the figures show (i) market equilibrium with three levels of β , for $\alpha = 1$ and (ii) market equilibrium with three levels of β , for $\alpha = 0$




Figure 7: The top figure shows market equilibrium *PPO* market shares for ranges of policies for α and β between 0 and 1, with full interactions. The middle figure shows market equilibrium δP for ranges of policies for α and β between 0 and 1, with full interactions. The bottom figure shows market equilibrium welfare outcomes for ranges of policies for α and β between 0 and 1, with full interactions.

A Appendix: Proofs

Proof of Proposition 1:

We consider a policy x that maintains the ordering of individuals' willingness-to-pay and thus the corresponding surplus from buying insurance conditional on the share of insured individuals $Q(x)$. We denote by $\tilde{w}(x)$ an individual's net willingness-to-pay and the corresponding density by $g^{\tilde{w}(x)}$. For the marginal individual, the net willingness-to-pay equals $\tilde{w}(x) = P(x) = D^{-1}(Q(x); x)$, while the original willingness-to-pay equals $\tilde{w}(0) = D^{-1}(Q(x); 0)$. Equilibrium welfare for policy x (not accounting for the budgetary cost) equals

$$\begin{aligned} \mathcal{W}(x) &= \int_{P(x)} E_{\tilde{w}(x)=\tilde{w}'}(s) g^{\tilde{w}(x)}(\tilde{w}') d\tilde{w}' \\ &= \int_{D^{-1}(Q(x); 0)} E_{w=w'}(s) g(w') dw'. \end{aligned}$$

The second equality follows by maintaining the ordering for any intensity of the policy x . Hence, the welfare effect of an increase in x only depends its impact on the marginal buyer. Using Leibniz' rule and $\frac{\partial}{\partial Q} D^{-1} = \frac{-1}{g^{\tilde{w}(x)}(P(x))}$, we find

$$\mathcal{W}'(x) = E_{P(x)}(s) Q'(x) = [P(x) - E_{P(x)}(c) - E_{P(x)}(f)] Q'(x).$$

Hence, a policy that increases Q , *ceteris paribus*, increases welfare if and only if $P(x) - E_{P(x)}(c) \geq E_{P(x)}(f)$. \square

Proof of Proposition 2:

The equilibrium price is characterized by $P^c(\alpha) = E_{\geq P^c(\alpha)}(c)$, where

$$E_{\geq P^c}(c) = \frac{1}{1 - G^{\tilde{w}(\alpha)}(P^c)} \int_{P^c}^{\infty} E(c|\tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}',$$

denoting an individual's net willingness-to-pay by $\tilde{w}(\alpha)$ and density by $g^{\tilde{w}(\alpha)}$. The corresponding equilibrium quantity equals $Q^c(\alpha) = D(P^c(\alpha); \alpha)$. We can solve for the impact of the policy on the equilibrium coverage by implicit differentiation of the equilibrium condition.

We consider first the impact of a change in the price on the equilibrium condition. Using Leibniz' rule, we find

$$\frac{\partial}{\partial P} \left[\int_{P^c} E(c|\tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}' \right] = -E_{P^c}(c) g^{\tilde{w}(\alpha)}(P^c).$$

Hence, using $\partial [1 - G^{\tilde{w}(\alpha)}(P)] / \partial P = -g^{\tilde{w}(\alpha)}(P)$, we find

$$\frac{\partial}{\partial P} E_{\geq P^c}(c) = [E_{\geq P^c}(c) - E_{P^c}(c)] \frac{g^{\tilde{w}(\alpha)}(P^c)}{1 - G^{\tilde{w}(\alpha)}(P^c)},$$

which depends on the difference between average and marginal cost.

We now consider the impact of a change in the policy on the equilibrium condition. We will show that

$$\frac{\partial}{\partial \alpha} \left[\int_{P^c} E(c|\tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}' \right] = -E_{P^c}(c \times f) g^{\tilde{w}(\alpha)}(P^c), \quad (5)$$

and, in an analogue way, we can find

$$\frac{\partial}{\partial \alpha} \left[\int_{P^c} g^{\tilde{w}(\alpha)}(w') dw' \right] = -E_{P^c}(f) g^{\tilde{w}(\alpha)}(P^c).$$

Hence,

$$\begin{aligned} \frac{\partial}{\partial \alpha} E_{\geq P^c}(c) &= [E_{\geq P^c}(c) E_{P^c}(f) - E_{P^c}(c \times f)] \frac{g^{\tilde{w}(\alpha)}(P^c)}{1 - G^{\tilde{w}(\alpha)}(P^c)} \\ &= [(E_{\geq P^c}(c) - E_{P^c}(c)) E_{P^c}(f) - cov_{P^c}(c, f)] \frac{g^{\tilde{w}(\alpha)}(P^c)}{1 - G^{\tilde{w}(\alpha)}(P^c)}, \end{aligned}$$

which again depends on the difference between average and marginal costs, but also on the covariance between costs and frictions among the marginal buyers, $cov_{P^c}(c, f)$.

The key step is to prove condition (5). Evaluating the impact of the information policy on $E_{\geq P^c}(c)$ is more involved than evaluating the impact of a price change, since the policy changes the sorting into insurance. However, we can re-write

$$\begin{aligned} & \int_{P^c} E(c | \tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}' \\ &= \int \int_{P^c} E(c | \tilde{w}(\alpha) = \tilde{w}', f = f') g^{\tilde{w}(\alpha)|f}(\tilde{w}'|f') d\tilde{w}' g^f(f') df' \\ &= \int \int_{P^c} E(c | w = \tilde{w}' + \alpha f', f = f') g^{w|f}(\tilde{w}' + \alpha f'|f') d\tilde{w}' g^f(f') df' \\ &= \int \int_{P^c + \alpha f'} E(c | w = w', f = f') g^{w|f}(w'|f') dw' g^f(f') df'. \end{aligned}$$

The first equality follows from the law of iterated expectations. The second equality follows from the identity $\tilde{w}(\alpha) = w - \alpha f$. Hence, conditional on $f = f'$, the density of the net willingness-to-pay $\tilde{w}(\alpha)$ equals the density of the willingness to pay w shifted by $\alpha f'$, $g^{\tilde{w}(\alpha)|f}(\tilde{w}'|f') = g^{w|f}(\tilde{w}' + \alpha f'|f')$. Moreover, conditioning on $\tilde{w}(\alpha) = \tilde{w}'$ and $f = f'$ is equivalent to conditioning on $w = \tilde{w}' + \alpha f'$ and $f = f'$. The last equality uses the substitution $w' = \tilde{w}' + \alpha f'$, for each $f = f'$, and thus $dw' = d\tilde{w}'$.

After this manipulation, we can simply apply Leibniz' rule again

$$\begin{aligned} & \frac{\partial}{\partial \alpha} \left[\int_{P^c} E(c | \tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}' \right] \\ &= \int \frac{\partial}{\partial \alpha} \left[\int_{P^c + \alpha f} E(c | w = w', f = f') g^{w|f}(w'|f') dw' \right] g^f(f') df' \\ &= - \int \left[E(c | w = P^c + \alpha f', f = f') g^{w|f}(P^c + \alpha f'|f') \times f' \right] g^f(f') df' \\ &= - \int \left[E(c \times f | w = P^c + \alpha f', f = f') g^{w|f}(P^c + \alpha f'|f') \right] g^f(f') df'. \end{aligned}$$

Using again the law of iterated expectations and the identity $\tilde{w}(\alpha) = w - \alpha f$, we thus indeed find

$$\frac{\partial}{\partial \alpha} \left[\int_{P^c} E(c | \tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}' \right] = -E_{P^c}(c \times f) g^{\tilde{w}(\alpha)}(P^c).$$

By implicit differentiation of the equilibrium condition $D^{-1}(Q^c(\alpha); \alpha) = E_{\geq D^{-1}(Q^c(\alpha); \alpha)}(c)$, expressed in quantities, we find

$$Q^c(\alpha) = -\frac{\left[1 - \frac{\partial}{\partial P^c} E_{\geq P^c}(c)\right] \frac{\partial D^{-1}(Q^c; \alpha)}{\partial \alpha} - \frac{\partial}{\partial \alpha} E_{\geq P^c}(c)}{\left[1 - \frac{\partial}{\partial P^c} E_{\geq P^c}(c)\right] \frac{\partial D^{-1}(Q^c; \alpha)}{\partial Q}}.$$

Using $\frac{\partial D^{-1}(Q^c; \alpha)}{\partial Q} = \frac{1}{-g^{\tilde{w}(\alpha)}(P^c)}$, $\frac{\partial D^{-1}(Q^c; \alpha)}{\partial \alpha} = -E_{P^c}(f)$, $\frac{g^{\tilde{w}(\alpha)}(P^c)}{1 - G^{\tilde{w}(\alpha)}(P^c)} = \frac{|\varepsilon_D(P^c)|}{P^c}$ and the expressions derived before, we find

$$Q^c(\alpha) = -\frac{E_{P^c}(f) - \text{cov}_{P^c}(c, f) \frac{|\varepsilon_D(P^c)|}{P^c}}{1 - [E_{\geq P^c}(c) - E_{P^c}(c)] \frac{|\varepsilon_D(P^c)|}{P^c}} g^{\tilde{w}(\alpha)}(P^c),$$

where the terms with $[E_{\geq P^c}(c) - E_{P^c}(c)]$ dropped out of the numerator. Finally, for a uniform subsidy S such that $P^c = E_{\geq P^c}(c) + S$, we have

$$Q^c(S) = \frac{1}{1 - [E_{\geq P^c}(c) - E_{P^c}(c)] \frac{|\varepsilon_D(P^c)|}{P^c}} g^{\tilde{w}(\alpha)}(P^c).$$

Defining this quantity effect as η^c , the first expression in the Proposition follows.

We now consider the impact on welfare. Welfare equals

$$\mathcal{W}(\alpha) = \int_{P(\alpha)} E(s | \tilde{w}(\alpha) = \tilde{w}') g^{\tilde{w}(\alpha)}(\tilde{w}') d\tilde{w}',$$

where $P(\alpha) = D^{-1}(Q(\alpha), \alpha)$. The total impact of the policy on welfare depends on the policy's effect on the equilibrium quantity and its direct effect on welfare,

$$\mathcal{W}'(\alpha) = \frac{\partial \mathcal{W}}{\partial P} \frac{\partial D^{-1}(Q(\alpha), \alpha)}{\partial Q} Q'(\alpha) + \frac{\partial \mathcal{W}}{\partial P} \frac{\partial D^{-1}(Q(\alpha), \alpha)}{\partial \alpha} + \frac{\partial \mathcal{W}}{\partial \alpha}.$$

By analogy to the above argument for the positive impact, we find

$$\begin{aligned} \frac{\partial \mathcal{W}}{\partial P} &= -E_{P(\alpha)}(s) g^{\tilde{w}(\alpha)}(P(\alpha)) \\ \frac{\partial \mathcal{W}}{\partial \alpha} &= -E_{P(\alpha)}(s \times f) g^{\tilde{w}(\alpha)}(P(\alpha)). \end{aligned}$$

Using $\frac{\partial D^{-1}(Q(\alpha); \alpha)}{\partial Q} = \frac{1}{-g^{\tilde{w}(\alpha)}(P(\alpha))}$ and $\frac{\partial D^{-1}(Q(\alpha); \alpha)}{\partial \alpha} = -E_{P(\alpha)}(f)$, we find

$$\begin{aligned} \mathcal{W}'(\alpha) &= E_{P(\alpha)}(s) Q'(\alpha) + E_{P(\alpha)}(s) E_{P(\alpha)}(f) g^{\tilde{w}(\alpha)}(P(\alpha)) - E_{P(\alpha)}(s \times f) g^{\tilde{w}(\alpha)}(P(\alpha)) \\ &= E_{P(\alpha)}(s) Q'(\alpha) - \text{cov}_{P(\alpha)}(s, f) g^{\tilde{w}(\alpha)}(P(\alpha)) \end{aligned}$$

and the second expression of the Proposition immediately follows as well. \square

Proof of Proposition 3 For a risk-adjustment policy β the competitive equilibrium is determined by

$$P^c(\beta) = E_{\geq P^c(\beta)}(\tilde{c}(\beta))$$

and $Q^c(\beta) = D^{-1}(P^c(\beta))$. By implicit differentiation, we find

$$Q^{c'}(\beta) = -\frac{-\frac{\partial}{\partial\beta}E_{\geq P^c}(\tilde{c}(\beta))}{\left[1 - \frac{\partial}{\partial P^c}E_{\geq P^c}(\tilde{c}(\beta))\right] \frac{\partial D^{-1}(Q^c;\beta)}{\partial Q}}.$$

Like in the proof of Proposition 1, we find $\frac{\partial D^{-1}(Q^c;\beta)}{\partial Q} = \frac{1}{-g(P^c)}$ and

$$\frac{\partial}{\partial P^c}E_{\geq P^c}(\tilde{c}(\beta)) = [E_{\geq P^c}(\tilde{c}(\beta)) - E_{P^c}(\tilde{c}(\beta))] \frac{g(P^c)}{1 - G(P^c)}.$$

Moreover,

$$\begin{aligned} \frac{\partial}{\partial\beta}E_{\geq P^c}(\tilde{c}(\beta)) &= \frac{\partial}{\partial\beta} \left[\frac{1}{1 - G(P^c)} \int_{P^c}^{\infty} E_{w'}(\tilde{c}(\beta)) g(w') dw' \right] \\ &= \frac{1}{1 - G(P^c)} \int_{P^c}^{\infty} E_{w'}(-[c - Ec]) g(w') dw' \\ &= E_{\geq P^c}(c) - Ec \end{aligned}$$

Hence,

$$\begin{aligned} Q^{c'}(\beta) &= \frac{E_{\geq P^c}(c) - Ec}{1 - [E_{\geq P^c}(c) - E_{P^c}(c)] \frac{|\varepsilon_D(P^c)|}{P^c}} g(P^c) \\ &= \eta^c \times [E_{\geq P^c}(c) - Ec], \end{aligned}$$

where η^c equals the equilibrium impact of a uniform subsidy.

Welfare equals

$$\mathcal{W}(\beta) = \int_{D^{-1}(Q^c(\beta))} E_w(s) g(w') dw'.$$

Hence,

$$\mathcal{W}'(\beta) = -E_{P^c}(s) \frac{1}{-g(P^c)} g(P^c) Q^{c'}(\beta).$$

This proves the second expression of the Proposition. \square

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 4.²⁴ 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.²⁵ We predict this distribution in a sophisticated manner that incorporates (i) past diagnostic information (ICD-9 codes) (ii) the Johns Hopkins ACG predictive medical software package (iii) a non-parametric model linking modeled health risk to total medical expenditures using observed cost data and (iv) a detailed division of medical claims and health plan characteristics to precisely map total medical expenditures to out-of-pocket expenses. The level of precision we gain from the cost model leads to more credible estimates of the choice parameters of primary interest (e.g., risk preferences and information friction impacts).

In order to most precisely predict expenses, we categorize the universe of total medical claims into four mutually exclusive and exhaustive subdivisions of claims using the claims data. These categories are (i) hospital and physician (ii) pharmacy (iii) mental health and (iv) physician office visit. We divide claims into these four specific categories so that we can accurately characterize the plan-specific mappings from total claims to out-of-pocket expenditures since each of these categories maps to out-of-pocket expenditures in a different manner. We denote this four dimensional vector of claims C_{it} and any given element of that vector $C_{d,it}$ where $d \in D$ represents one of the four categories and i denotes an individual (employee or dependent). After describing how we predict this vector of claims for a given individual, we return to the question of how we determine out-of-pocket expenditures in plan j given C_{it} .

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

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

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

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

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

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

²⁴The model is similar to that used in Handel (2013).

²⁵In 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 $G_{I_{dz}}(\cdot)$. Then we assume that each individual in this cell has a distribution equal to a continuous fit of $G_{I_{dz}}(\cdot)$, which we denote G_{dz} :

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

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

The above process generates a distribution of claims for each d and z but does not model 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') .²⁶ Here, G_{idt} is the distribution G_{dz} where $i \in I_{dz}$ at time t . Since correlations are modeled across d we pick the metric θ to group people into cells for the basis of determining correlations (we use the same cells that we use to determine group people for hospital and physician office visit claims). Denote these cells based on θ by z_θ . Then for each cell z_θ denote the empirical rank correlation between claims of type d and type d' by $\rho_{z_\theta}(d, d')$. Then, for a given individual i we determine the joint distribution of claims across D for year t , denoted $H_{it}(\cdot)$, by combining i 's marginal distributions for all d at t using $\rho_{z_\theta}(d, d')$:

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

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

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

$$\Omega_j : C \rightarrow OOP_j$$

This mapping takes a given draw of claims from H_{it} and converts it into the out of pocket expenditures an individual would have for those claims in plan j . This mapping accounts for plan-specific 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

²⁶It 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 than 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:

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

Cost Model Estimation Structure

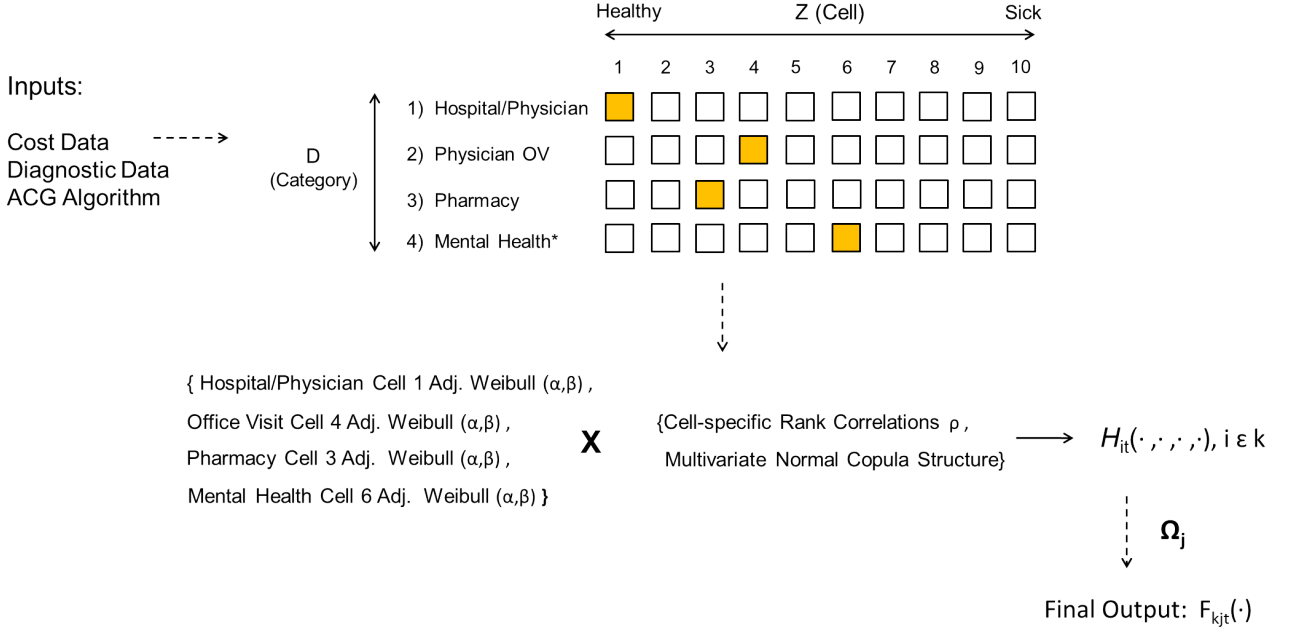


Figure B1: This figure outlines the primary steps of the cost model described in Appendix B. It moves from the initial inputs of cost data, diagnostic data, and the ACG algorithm to the final output F_{kjt} which is the family, plan, time specific distribution of out-of-pocket expenditures that enters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

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

Again, we use the notation $G_{iDt}^{\hat{}}$ 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 $G_{iDt}^{\hat{}}$ 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 $G_{iDt}^{\hat{}}$ 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 $\Phi_{1|2|3|4}^i$ 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:

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

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 Identification and 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 t_4 . 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 3:²⁷

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

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):²⁸

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

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 3. 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:

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

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

²⁷While we discuss estimation for the full model, the logic extends easily to the other specifications estimated in this paper.

²⁸Here, 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:

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

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:

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

Here d_{kj} 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.

A1 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 F, 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 F. 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 F 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: Simulations

To go beyond the local evaluations and provide further insights on how the different model components impact positive and normative outcomes under different policies, we present a series of simulations. We use these simulations to illustrate the role that the key micro-foundations described in this section play in determining market outcomes under (i) no policy interventions (ii) friction-reducing policies and (iii) risk-adjustment policies. Specifically, we distinguish between cases where friction-reducing interventions have positive vs. negative welfare impacts, and cases where effective risk-adjustment policies are essential prior to implementing friction-reducing policies.

Our focus is on a market setup in the mold of Einav et al. (2010), similar to our primary model, where insurers compete to offer supplemental insurance relative to a baseline publicly provided plan. See Appendix E for similar simulations on markets with two competitively priced plans, as studied in Handel et al. (2015).

The baseline plan for these simulations has a deductible of \$3,000, with 10% coinsurance after that point, up to an out-of-pocket maximum of \$7,000 (this plan has a 66% actuarial value for our baseline costs below). The supplemental coverage that insurers compete to offer covers all out-of-pocket spending in the baseline plan, and thus brings all consumers up to full insurance. These plans are similar to the minimum and maximum coverage levels regulated in the state-based exchanges set up in the ACA, and also mimic the plans we study in our empirical environment later in this paper. Importantly, in our environment with risk averse consumers and no moral hazard, all consumers purchase full insurance in the first-best. In each simulation we simulate the market for 10,000 consumers.

We study a range of scenarios that vary in terms of the underlying means and variances of (i) consumer surplus from risk protection (ii) consumer costs and (iii) consumer choice frictions. Table D1 describes the underlying distributions for the different cases we study. We simulate two scenarios for consumer yearly expected costs: both have the same mean of just above \$5,000. The first scenario has a high standard deviation of expected costs in the population of \$6,819 while the second has a low standard deviation of \$2,990. For each scenario, consumer expected costs are drawn from a lognormal distribution. The within-year standard deviation in costs for a given consumer is 3,000 plus 1.2 times their yearly expected costs in both scenarios, with each consumer's costs drawn from lognormal distributed as well. The impact of frictions on demand for generous insurance is generated from a normal distribution. The high (low) mean is a \$2,500 (\$0) shift in willingness-to-pay while the high (low) standard deviation we study is \$2,000 (\$500). We study all four combinations of these high/low means and variances. Finally, for consumer risk aversion, we also study four combinations from normal distributions with high/low means and variances. The high (low) CARA mean is $1 * 10^{-3}$ ($4 * 10^{-4}$) while the high (low) standard deviation is $4 * 10^{-4}$ ($1 * 10^{-4}$), with values truncated above 0.²⁹ The left panel in Figure 3 shows the two distributions of costs studied. The right panel in Figure 3 shows the distribution of surplus in the market when the variance in costs is high under the cases of (i) high mean and variance of risk aversion (ii) low mean and high variance of risk aversion and (iii) low mean and low variance of risk aversion.

We first present a specific simulation example to illustrate the very different impact frictions can have on equilibrium and welfare depending on the primitives of the model. We then systematically investigate positive and normative patterns across a wider range of simulations. The example we start with focuses on two markets that differ only in terms of mean surplus: one has low mean surplus and the other high mean surplus. Otherwise, both markets have a high variance of costs and surplus, and a low mean, but high variance of frictions.

²⁹The variance in surplus is likely to increase further when allowing for horizontally differentiated plans.

Simulations

Key Micro-Foundations

Costs - μ_c^*	Total costs - mean	5,373
Costs - High σ_c^*	Total costs - sd	6,819
Costs - Low σ_c^*	Total costs - sd	2,990
Surplus - High μ_s^{**}	CARA - mean	$1 * 10^{-3}$
Surplus - Low μ_s^{**}	CARA - mean	$3 * 10^{-4}$
Surplus - High σ_s^{**}	CARA - sd	$4 * 10^{-4}$
Surplus - Low σ_s^{**}	CARA - sd	$1 * 10^{-4}$
Frictions - High μ_f^{***}	WTP shift - mean	2,500
Frictions - Low μ_f^{***}	WTP shift - mean	0
Frictions - High σ_f^{***}	WTP shift - sd	2,000
Frictions - Low σ_f^{***}	WTP shift - sd	500

*Total costs simulated from lognormal distribution.

**CARA risk preferences simulated from normal distribution, truncated above 0.

***Shift in relative willingness-to-pay for high-coverage contract simulated from normal distribution.

Table D1: This table presents the parameter values of the distributions underlying the micro-foundations in our model for the different simulation scenarios we study.

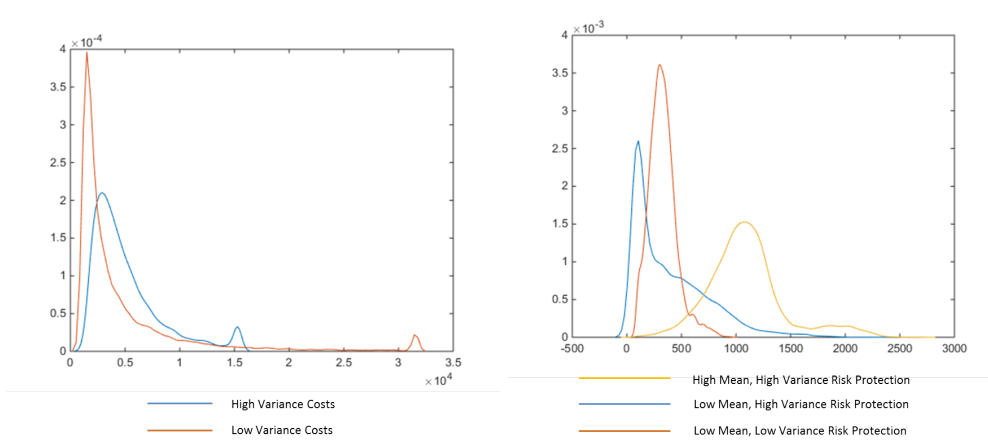


Figure D2: The left panel shows the two different distributions of total costs used in our simulations. The right panel shows the resulting surplus distributions under the different scenarios for the distribution of risk preferences, conditional on the cost distribution with high variance.




Figure D3: This figure shows the key market micro-foundations for the market with low mean surplus μ_s , in addition to high σ_s , low μ_f , high σ_f , and high σ_c . From left to right, the figure shows the three cases of (i) full frictions (ii) half frictions and (iii) no frictions.




Figure D4: This figure shows the average cost curves, as a function of how much frictions are reduced for the specific simulation example with high σ_s , high σ_f , high σ_c . The mean surplus and friction do not affect this figure as they maintain the ordering of consumers.

Figure 4 shows the key micro-foundations of the market with low mean surplus for the three policy cases of full frictions, frictions reduced by 50%, and no frictions. The figure illustrates a number of properties of markets with low surplus relative to costs when the variance of frictions is meaningful. When full frictions are present, the demand curve is more heavily skewed due to impacts of very positive and negative friction draws. The variance in frictions swamps the variance in costs and surplus, and the market holds together, with quantity of incremental coverage purchased equal to 0.51. When frictions are reduced by 50% ($\alpha = 0.5$) the variation in willingness-to-pay becomes much closer to the variation in costs and value, but the presence of frictions still helps hold the market together, with quantity of incremental coverage equal to 0.41. When frictions are fully removed, the market almost completely unravels, with only 11% of consumers buying incremental coverage. As the figures reveal, as frictions are reduced in this environment, the demand curve becomes less skewed, making it harder to hold the market together at the top end. Consumers with the highest willingness to pay tend to overestimate the insurance value the most and the friction reducing policy reduces their demand for insurance. In addition, as Figure 5 shows, the average cost curves become steeper as frictions are reduced, reflecting increased sorting based on costs.

Figure 6 shows the key micro-foundations of the market with high mean surplus for the same policy interventions. In contrast to the market with low surplus, this case illustrates properties of markets where friction-reducing policies can be beneficial. In this case, when full frictions are present, 64% of consumers purchase coverage in equilibrium. Now, however, when frictions are reduced by 50%, the equilibrium percentage purchasing coverage increases to 79%, and when no




Figure D5: This figure shows the key market micro-foundations for the market with high mean surplus μ_s , in addition to high σ_s , low μ_f , high σ_f , and high σ_c . From left to right, the figure shows the three cases of (i) full frictions (ii) half frictions and (iii) no frictions.

frictions are present the percentage with coverage increases further to 91%. Here, friction-reducing policies have a positive impact on equilibrium coverage. As discussed before, when the share of consumers purchasing coverage is high (due to the high mean surplus), the marginal consumers are more likely to have a bias against purchasing more coverage (i.e., the marginal friction value is negative). As frictions are reduced, these consumers have that bias reduced so that the demand for insurance increases. The level effect of the policy is thus positive in this market. The incremental sorting based on costs when frictions are reduced is the same as in the market with low mean surplus, but this sorting effect is now more than offset by the reverse level effect so that equilibrium coverage increases. Note that the policy not only increases equilibrium coverage, but also increases the match quality and thus will improve welfare as well (as discussed shortly).

We now investigate a broad range of scenarios corresponding to different combinations of the underlying market micro-foundations. Table D2 shows the proportion of consumers buying supplemental insurance as a function of these different micro-foundations. We explore comparative statics for different cases with full frictions present and investigate what happens when those frictions are reduced.

There are several notable patterns. First, conditional on the population distributions of surplus from risk protection s and costs c , reducing the mean level of frictions (which favor purchasing generous coverage) reduces the overall demand for insurance and thus unambiguously reduces the equilibrium quantity purchased. More interestingly, following Proposition ??, it is clear that the equilibrium implications of reducing the variance in frictions very much depends on the variance of costs. Comparing the first and second columns to the third and fourth columns shows that whether σ_c is high or low has an important impact on the degree of market unraveling as the mean and variance of frictions are reduced. For example, fixing μ_s as low and σ_s as high, when the frictions mean and variance is high, the market share of equilibrium coverage is 0.92 with low σ_c and 0.91 with higher σ_c . When the frictions changes to low mean, high variance, these quantities are 0.56 and 0.51 respectively. But, when the variance in frictions is also reduced to low (along with the mean), these quantities are 0.53 and 0.17. When frictions are fully removed, 39% of consumers purchase more generous coverage for this low σ_c case, but only 11% do in this high σ_c case. Thus, when potential surplus in the market is relatively low, high σ_c implies that reducing market frictions could be especially damaging for market function.

Table D2 also confirms that the mean and variance of surplus (relative to costs and frictions) have important implications for whether frictions are ‘good’ or ‘bad’ for market function. In the cases with low σ_c , low μ_s , and low μ_f , moving from high friction variance to low friction variance has little impact on the equilibrium quantity. When μ_s and μ_f are low, but σ_c is high, moving from high to low σ_f facilitates substantial unraveling (e.g., 0.51 to 0.17 purchasing under high σ_s).

Simulations
Equilibrium Quantities

	Low σ_c	Low σ_c	High σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	Low μ_s	Low μ_s	High μ_s
	Low σ_s	High σ_s	Low σ_s	High σ_s	High σ_s
High μ_f , High σ_f	0.93	0.92	0.90	0.91	0.95
High μ_f , Low σ_f	1	1	1	1	1
Low μ_f , High σ_f	0.59	0.56	0.49	0.51	0.64
Low μ_f , Low σ_f	0.65	0.53	0.07	0.17	0.79
No Frictions	0.45	0.39	0.06	0.11	0.91

Table D2: This table presents the proportion of the market purchasing incremental insurance in equilibrium, for a range of underlying population micro-foundations.

However, even with high σ_c , when μ_s and σ_s are high, with low μ_f reducing the variance of frictions *increases* the equilibrium quantity from 0.64 to 0.79. The role (negative) frictions play in pushing people away from generous coverage outweighs the role that they play in reducing adverse selection through sorting.

Taken all together, these results support the earlier analysis by illustrating that (i) reducing the mean impact of frictions on willingness-to-pay for insurance always reduces insurance coverage (ii) reducing the variance and impact of frictions can be good when the mean surplus is relatively high, but (iii) incremental adverse selection occurs and reduces coverage more when the variance in costs is relatively high. These results are further borne out in the bottom of Table D2, which studies the same scenarios, but under the policy where frictions are completely eliminated ($\alpha = 1$). In Appendix F, in Table F5, we also present results for simulations for $\alpha = 0.5$, or partially-reduced frictions, with the comparative statics intuitively following the patterns already described here.

Table D3 presents the proportion of the first-best surplus achieved in each scenario. Notably, welfare is increasing for friction-reducing policies when μ_s and σ_s are high, but decreasing when those values are lower. The sensitivity of the relationship to the level of σ_c is substantial: when σ_c is low the market does not unravel when frictions are reduced, but when σ_c is high it unravels rather quickly and so does the surplus achieved. The welfare implications tend to be in line with the implications for market function: equilibrium surplus increases when equilibrium coverage increases and vice-versa. The exception holds when the variance of surplus is high relative to the variance of costs. In particular, moving from high σ_f to low σ_f (keeping μ_f low), we find that equilibrium coverage decreases, while equilibrium surplus increases. The reason is that the positive matching effect of reduced frictions outweighs the negative equilibrium consequences of any incremental selection on costs, in line with the trade-off highlighted in Proposition ???. Table F4 in Appendix F illustrates the improved matching by showing how the proportion of mistakes consumers make reduces, given the equilibrium price in each scenario. This table also highlights that while the market unraveling due to friction-reducing policies may decrease total welfare, some consumers will be better off as they now avoid making mistakes.³⁰

Both the magnitude and direction of the welfare impact that friction-reducing policies have depend on how effective risk-adjustment transfers in the market are in mitigating adverse selection. Table D4 studies the interaction between friction-reducing policies and risk-adjustment policies.

³⁰This table shows, among other things, that the proportion of mistakes made in equilibrium is only related to surplus achieved in the market when the mean and variance of surplus are large relative to frictions and costs.

Simulations**Equilibrium Surplus**

	Low σ_c	Low σ_c	High σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	Low μ_s	Low μ_s	High μ_s
	Low σ_s	High σ_s	Low σ_s	High σ_s	High σ_s
High μ_f , High σ_f	0.94	0.94	0.90	0.94	0.95
High μ_f , Low σ_f	1	1	1	1	1
Low μ_f , High σ_f	0.62	0.61	0.51	0.61	0.67
Low μ_f , Low σ_f	0.72	0.66	0.14	0.33	0.84
No Frictions	0.56	0.55	0.10	0.23	0.95

Table D3: This table presents the proportion of first-best surplus achieved in the market for a range of underlying population micro-foundations.

We use the underlying distribution of frictions with low μ_f and high σ_f for all risk-adjustment scenarios.

Consider first the case with high σ_c , low μ_s , and high σ_s . When there is no risk-adjustment the market unravels and welfare decreases as frictions are reduced. With partially effective risk-adjustment ($\beta = 0.5$), reducing frictions still reduces equilibrium quantity and welfare, but by a much lesser degree. With full risk-adjustment ($\beta = 1$), there is almost no impact of reduced frictions on quantity, and welfare increases as frictions are reduced. Thus, in this scenario, friction-reducing policies become more tenable, and switch from 'bad' to 'good' as risk-adjustment is more effective. Figure 7 shows the outcomes in this market under full frictions and under no frictions for the three different risk-adjustment scenarios studied.

Compare this now to the case with low σ_c , keeping μ_s low, and σ_s high. With no risk-adjustment reducing frictions has a slight negative impact on equilibrium quantity and welfare. With partial risk-adjustment as frictions are reduced quantity is relatively unchanged but welfare increases substantially, reflecting the impact of better consumer-plan matches. Under full risk-adjustment, both quantity and welfare are strongly increasing as frictions are reduced. Finally, column 3 demonstrates that in the case of high mean surplus for which friction-reducing policies were good for equilibrium quantity and welfare even under no risk-adjustment, this gradient increases as risk-adjustment becomes more effective. Taken in sum, as the mean and variance of surplus increase relative to the mean and variance of costs in the population, the threshold of risk-adjustment necessary to make friction-reducing policies have a positive welfare impact is decreasing.

While our analysis focuses on the case where there is one type of supplemental insurance that is competitively provided, as in Einav et al. (2010), much of the intuition presented in this section extends to the type of market where two classes of plans with different actuarially levels are competitively offered (see e.g., Handel et al. (2015) or Weyl and Veiga (2017)). The key difference in practice between these two types of markets is that the market for supplemental coverage is less likely to unravel, because the supplemental insurer covers only incremental costs rather than the total costs of the sickest consumers. The comparative statics we study remain the same in spirit for this alternative market design: Appendix E presents simulation analysis similar to that presented in this section, but for the case of two priced classes of insurance offerings.³¹

³¹Several insights emerge. First, for a given set of micro-foundations, these multi-plan markets are much more likely to unravel. Consequently, the mean and variance of surplus relative to costs must be substantially higher for friction-reducing policies to have positive impacts in those markets, conditional on a given level of risk-adjustment. In markets with two priced plans, friction-reducing policies are always beneficial under full risk-adjustment, but risk-adjustment

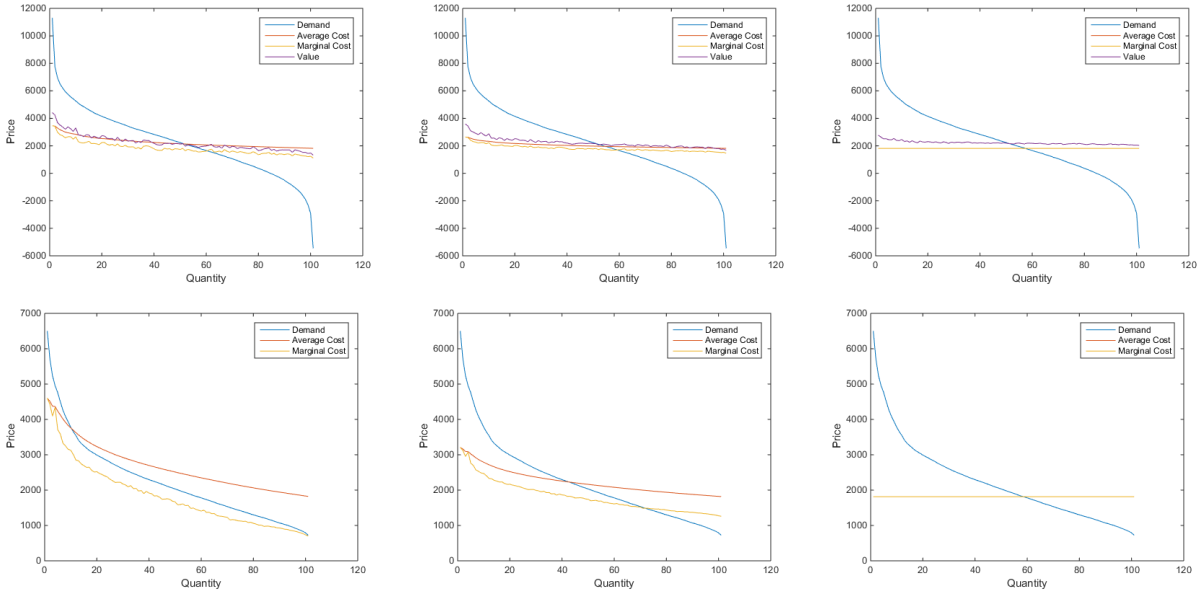


Figure D6: This figure shows market outcomes under different risk-adjustment transfer effectiveness levels. The top panel shows the impact of risk-adjustment with full frictions present, while the bottom shows the impact of risk-adjustment when no frictions are present. The market studied has high σ_c low μ_s , high σ_s , low μ_f , and high σ_f . From left to right, the figure shows the three cases of (i) no risk-adjustment (ii) partial risk-adjustment and (iii) full risk-adjustment.

must be much more effective than in the market for supplemental coverage to make friction-reducing policies welfare increasing. Thus, while the same basic intuition holds in markets with two plan types, policymakers should have a higher threshold for the effectiveness of risk-adjustment when considering the implementation of friction-reducing policies. See Appendix E for more detail on these markets, commonly referred to as exchanges.

Simulations			
Risk-Adjustment			
Quantity			
(% Surplus Achieved)			
	Low σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	High μ_s
	High σ_s	High σ_s	High σ_s
No Risk-Adjustment ($\beta = 0$)			
Full Frictions	0.56 (61%)	0.51 (61%)	0.64 (67%)
Half Frictions	0.57 (66%)	0.41 (56%)	0.72 (76%)
No Frictions	0.39 (55%)	0.11 (23%)	0.91 (95%)
Partial Risk-Adjustment ($\beta = .5$)			
Full Frictions	0.57 (62%)	0.54 (64%)	0.66 (70%)
Half Frictions	0.61 (69%)	0.52 (66%)	0.75 (79%)
No Frictions	0.62 (79%)	0.42 (60%)	0.93 (96%)
Full Risk-Adjustment ($\beta = 1$)			
Full Frictions	0.58 (64%)	0.57 (66%)	0.68 (72%)
Half Frictions	0.64 (72%)	0.59 (72%)	0.77 (81%)
No Frictions	0.71 (86%)	0.58 (75%)	0.94 (97%)

Table D4: This table presents equilibrium quantity sold, and proportion of total surplus achieved, as a function of the underlying risk-adjustment (β) and friction-reducing policies (α). The entire Table considers the case of low μ_f and high σ_f .

E Appendix: Equilibrium with Two Types of Competing Plans

The primary empirical analysis discussed in the text is for an insurance market where there is a basic government option provided and insurers compete to provide supplemental insurance. As noted in the text, this setup is in the spirit of Einav et al. (2010). An alternative setup we describe in Section 2 and in Section ?? is that where insurers compete to offer two types of insurance plans simultaneously (so costs for both types of plans must break even with premiums in equilibrium). This latter setup is in the spirit of recent work by Handel et al. (2015) studying equilibrium in insurance exchanges. In an example, Veiga and Weyl (2016) illustrate how these two types of setups can lead to markedly different results, primarily because when costs are endogenized for basic coverage the costs incurred by each plan are similar to total expected costs, while when coverage is supplemental, costs are similar only to incremental spending in the supplemental coverage. Thus, the costs faced by the insurers providing generous coverage in the Handel et al. (2015) setup are substantially larger than those when the coverage insurers compete to offer is supplemental. This makes it more likely that equilibrium will unravel towards less generous coverage, because incremental premiums for generous coverage must reflect this larger cost difference.

The analysis in Section 2 considered the choice to buy incremental insurance from a competitive market or stick with a baseline option. Our comparative statics for how key micro-foundations interact with friction-reducing and risk-adjustment policies (and how those foundations determine equilibrium in the absence of such policies) remain the same in the case of more than one type of priced plan. The primary change is that both the high and low coverage plans must account for sorting based on costs in premium setting, whereas in the supplemental insurance case there is no premium for baseline coverage, so it does not adjust along with endogenous sorting. In practice, as shown in Weyl and Veiga (2017) and Handel et al. (2015), this internalization of costs by both plan types leads to unraveling in the market that is of an order of magnitude higher, conditional on the same population consumer micro-foundations.

We briefly illustrate this for a choice between two plans, a low-coverage plan L providing only and a high-coverage plan H . If both plans are priced in competitive markets (as in Weyl and Veiga (2017) and Handel et al. (2015)), each plan needs to internalize the full cost of its own consumers.

We relate our six potentially relevant dimensions of heterogeneity to our original setup as follows:

$$c = c^H - c^L, s = s^H - s^L, f = f^H - f^L \text{ and } P = P^H - P^L.$$

Since for each plan type the price equals the average cost of the individuals selecting the respective plan, the price differential equals

$$\begin{aligned} P &= E_{\geq P}(c^H) - E_{< P}(c^L) \\ &= E_{\geq P}(c) - [E_{< P}(c^L) - E_{\geq P}(c^L)]. \end{aligned}$$

The second term captures the difference in baseline coverage costs between those actually buying the low-coverage plan relative to those buying the high-coverage plan. If the baseline coverage costs are independent of the sorting of individuals, the previous equilibrium analysis entirely generalizes. If not, we need to take into account how the policy affects the selection based on the full cost into both plans.³² Similarly, when evaluating policy interventions, the re-sorting based on the full cost in both plans determines the new equilibrium prices. The welfare analysis naturally generalizes

³²If both plans insure the same underlying risk but differ in their overall coverage, we have "adverse selection" into the high-coverage contract, both for the baseline and supplemental coverage, i.e., $E_{\geq P}(c^L) \geq E_{< P}(c^L)$. If two plans insure different types of risk, we may well have "adverse selection" into both contracts, i.e., $E_{< P}(c^L) \geq E_{\geq P}(c^L)$.

since the change in total welfare only depends on the change in the differential surplus (as long as the purchase of baseline coverage is mandated). That is,

$$\begin{aligned}\mathcal{W} &= (1 - G(P)) E_{\geq P}(s^H) + G(P) E_{< P}(s^L) \\ &= (1 - G(P)) E_{\geq P}(s) + E(s^L).\end{aligned}$$

See Handel et al. (2015) and Weyl and Veiga (2017) for a much more complete discussion of equilibrium in markets with multiple tiers of competitively priced plans, and how they compare to the market with baseline coverage and privately-provided supplemental coverage. For this paper, it is only important to note that the comparative statics will be the same directionally, regardless, though of course the threshold for what makes a market unravel vs. not it much lower in the markets with two or more types of priced plans.

In Section D we presented simulations to illustrate the relationship between market micro-foundations and different policy recommendations, in the Einav et al. (2010) style market with one priced supplemental plan. Here, in Table E1 we present analogous results for the market with two priced plans. The underlying simulation micro-foundations for each scenario are the same as those described in the main text in Table D1.

In the supplemental market described in Section D, for the scenarios where the distribution of surplus was high relative to costs, the equilibrium held together and friction-reducing policies were welfare improving. In the market for two priced plans, this is not the case. With high mean and variance of frictions the case with high mean and variance of surplus has quantity equal to 0.55. When frictions are reduced, either by 50% or 100% the market completely unravels, in contrast to the supplemental market. This is for the case where the variance in costs is high. For the other two scenarios presented in Table E1, with low variance in costs, low mean surplus, and low or high surplus variance, the results have a similar flavor across the range of frictions present and friction policies. With high mean frictions, the market holds together and quantity provided is high. But, when the mean level of frictions is low, or policies are in place to reduce the high mean frictions, the market fully unravels and no generous insurance is purchased in equilibrium.

These results imply that the market with two priced plans is much more likely to unravel for a given set of micro-foundations. As a result, in this style market, policies to reduce frictions are more likely to be welfare decreasing than in the market with one competitively priced supplemental plan. While the mean and variance of surplus relative to the mean and variance of costs is still a crucial determinant of whether friction-reducing policies will be good or bad, now because of the nature of the market the distribution of insurance must be higher relative to the distribution of costs in order for the market to function and in order for friction-reducing policies to be welfare positive.

A corresponding implication is that the threshold for risk-adjustment that is necessary to make friction-reducing policies welfare positive is higher in the market with two priced plans. Table E2 presents market quantities and welfares for a range of interacted risk-adjustment and friction-reducing policies. As in the main text, results are presented for the case with low μ_f and high σ_f .³³ It is clear that in all cases studied, incremental risk-adjustment increases welfare and is absolutely crucial when implementing friction-reducing policies in the market. For any of the cases presented, when risk-adjustment is either partially effective ($\beta = 0.5$) or not present ($\beta = 0$) friction-reducing policies reduce equilibrium coverage and increase adverse selection. However, when full risk-adjustment is present, friction-reducing policies improve equilibrium quantity and welfare in

³³Note that when the mean level of frictions are increased, the equilibrium is less likely to unravel, we present this case so it can be directly compared to the supplemental equilibrium in the text.

Simulations-Two Priced Plans
Equilibrium Quantities

	Low σ_c	Low σ_c	High σ_c
	Low μ_s	Low μ_s	High μ_s
	Low σ_s	High σ_s	High σ_s
Full Frictions			
High μ_f , High σ_f	0.83	0.79	0.55
High μ_f , Low σ_f	0.9	0	0
Low μ_f , High σ_f	0.33	0.28	0
Low μ_f , Low σ_f	0	0	0
No Frictions	0	0	0
Half Frictions			
High μ_f , High σ_f	0	0	0
High μ_f , Low σ_f	0	0	0
Low μ_f , High σ_f	0	0	0
Low μ_f , Low σ_f	0	0	0
No Frictions	0	0	0

Table E1: This table presents the proportion of the market purchasing full insurance in equilibrium, for a range of underlying population micro-foundations. These results are for an insurance exchange where two types of plans are offered competitively.

the market.

Thus, the same underlying intuition holds for markets with two priced plans, but the threshold for what constitutes ‘enough’ risk-adjustment to implement friction-reducing policies is much higher because of the higher potential for adverse selection. This distinction is generally interesting, and reflects the underlying notion that, as the mean and variance of population costs becomes high relative to the mean and variance of surplus from risk protection, friction-reducing policies are more likely to be welfare -decreasing and more risk-adjustment is required for them to be welfare-increasing.

In addition to presenting these simulations, we also conduct the analog to our empirical analysis in the text for the case of two priced plans. Section ?? lays out the model for insurer competition in both the Einav et al. (2010) and Handel et al. (2015) cases. Since the mean and variance of costs are high relative to surplus in our empirical application, it is highly likely that the market will unravel except for cases with very high frictions or very effective risk-adjustment.

Figure E shows market equilibrium for the baseline case where $\alpha = 0$ and $\beta = 0$. The average cost line for generous coverage always lies above the demand curve, even in this case where substantial mean frictions push people towards that coverage. The high mean and variance of consumer costs, relative to their surplus from incremental coverage, leads to this scenario, which we also see in the simulation in Section 2. Not surprisingly, when frictions are partially and fully removed in figures E7 and E8 there is still no positive equilibrium market share of more generous coverage, which is expected given that the frictions we estimate push consumers toward that coverage.

There is some hope for maintaining generous coverage when there is insurer risk-adjustment.




Figure E7: Market Equilibrium Including Information Frictions




Figure E8: Market Equilibrium with Partial Information Frictions




Figure E9: Market Equilibrium without Information Frictions

Simulations-Two Priced Plans			
Risk-Adjustment			
Quantity			
(% Surplus Achieved)			
	Low σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	High μ_s
	High σ_s	High σ_s	High σ_s
No Risk-Adjustment ($\beta = 0$)			
Full Frictions	0.28 (35%)	0 (0%)	0 (0%)
Half Frictions	0 (0%)	0 (0%)	0 (0%)
No Frictions	0 (0%)	0 (0%)	0 (0%)
Partial Risk-Adjustment ($\beta = .5$)			
Full Frictions	0.45 (50%)	0.01 (2%)	0.25 (30%)
Half Frictions	0.07 (11%)	0 (0%)	0 (0%)
No Frictions	0 (0%)	0 (0%)	0 (0%)
Full Risk-Adjustment ($\beta = 1$)			
Full Frictions	0.58 (64%)	0.57 (66%)	0.68 (72%)
Half Frictions	0.64 (72%)	0.59 (72%)	0.77 (81%)
No Frictions	0.71 (86%)	0.58 (75%)	0.94 (97%)

Table E2: This table presents equilibrium quantity sold, and proportion of total surplus achieved, as a function of the underlying risk-adjustment (β) and friction-reducing policies (α). The entire table considers the case of low μ_f and high σ_f . Results presented are for the market with two competitively priced plans.

Figure E shows that some coverage is possible with full frictions and with either partial or full risk-adjustment. When frictions are removed, even with full risk-adjustment there is full unraveling of the market: this is because the cost of the average consumer for the family tier we study is higher than the top-end value of insurance coverage, given the way that the insurance contracts are set up relative to one another. Handel et al. (2015) shows that equilibrium in the market is harder to maintain the closer the two types of coverage are relative to one another, precisely for this reason.

Thus, with the limited surplus estimated in our environment from risk-protection, and the closeness of the two types of insurance contracts relative to average costs, the market outcome in our environment is almost always full unraveling. There are a few reasons why we might not see this in practice. First, consumers typically receive subsidies to purchase insurance coverage, either from the government in exchanges or from their employer in employer provided insurance. Though one typical principle of managed competition is that consumers receive a lump sum subsidy and pay the full marginal cost of generous coverage, in practice in many exchanges poorer consumers have caps on the premiums that they pay, limiting the relative premium spread between insurance contracts. The second reason is that consumers with frictions may follow decision models whereby they always choose more generous coverage no matter what. With the micro-foundations in our




Figure E10: Market Equilibrium with full frictions and a range of risk-adjustment policies.




Figure E11: Market Equilibrium with no frictions and a range of risk-adjustment policies.

environment, even the presence of such consumers would not hold the equilibrium together, given the spread because average costs in the PPO and the relative generosity of that coverage, unless the consumers choosing generous coverage by mistake are the healthiest in the population.

F Appendix: Additional Analysis

Table E1 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 the four years of data (t_1 - t_4) that we observe.³⁴ 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 t_3 , and use the re-weighted statistics for the second and third columns, as described in the text.

Table F2 presents the details of plan design for the two plans consumers choose between in our empirical environment. Table F5 presents the results for the simulations in Section 2.5 that are for the case of partially effective friction-reducing policies ($\alpha = 0.5$). Table F4 describes the proportion of consumers making choice mistakes in each of the simulation scenarios described in Section 2.5.

Figure E1 presents the smoothed distribution of expected costs from the cost model for families in the primary sample.

Figure F1 depicts the financial returns to selecting the HDHP option relative to the PPO option for an employee in the family tier, which has more than 50% of the employees in our sample.³⁵ 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. For a family, the range of potential ex-post value for the HDHP spans $[-\$2,500, +\$3,750]$, with the lower bound coming from cases with a lot of medical spending, the upper bound coming from the case of zero spending.³⁶ Based on ex post spending 60% of employees, across all tiers, are better off financially in the HDHP, though only 15% of employees actually choose that plan.

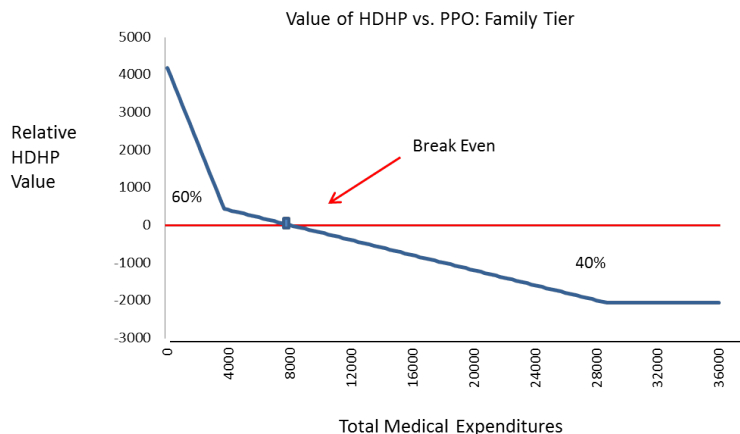


Figure F1: Description of HDHP financial value relative to the PPO in year t_4 , for the family tier, as a function of total medical expenditures. This chart assumes that employees contribute 50% of the maximum possible incremental amount to their HSA, near the median in the population. 60% of *all* employees would be better off ex-post in the HDHP, given their respective coverage tiers.

Table F3 presents the results from the primary choice model we use in the main text, discussed

³⁴We cannot provide the exact number of overall employees, to preserve the anonymity of the firm. As noted earlier, we cannot state the exact years of the data, though we note they are from a four-year period between 2008-2014.

³⁵The same general structure holds for couples and families with shifts in the levels of the key plan terms.

³⁶This range shifts upward by a constant amount if consumers derive value from incremental HSA contributions: the figure assumes consumers contribute 50% of the potential incremental contribution, up to the maximum allowed, consistent with what we find in the data. Note also that the relative value range for an individual / couple equals the family bounds multiplied by 0.4 (0.8).

Sample Demographics			
	Full Sample	Survey Recip. (Weighted)	Survey Resp. (Weighted)
N -Employees	[35,000 , 60,000]*	4500	1661
N_d - Emp.& Dep.	[105,000 , 200,000]*	11,690	4,584
t_3 PPO%	88.8	89.6	88.7
t_4 PPO%	82.7	83.0	81.6
t_3 HDHP %	11.2	10.4	11.3
t_4 HDHP %	17.3	17.0	18.4
Gender, Emp. and Dep. (% Male)	51.8	51.5	51.1
Age			
18-29	8.6%	14.9%	11.6%
30-39	41.1%	43.8%	42.7%
40-49	38.1%	32.7%	34.1%
50-59	10.9%	7.7%	10.5%
≥ 60	1.3%	0.9%	1.2%
Income			
Tier 1 (< \$100K)	12.8%	15.3%	16.2%
Tier 2 (\$100K-\$150K)	65.8%	68.5%	69.2%
Tier 3 (\$150K-\$200K)	16.7%	14.3%	12.9%
Tier 4 (\$200K+)	3.5%	1.2%	0.9%
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 F1: This table gives summary statistics for the employees and dependents of the firm we use data from. When not stated, statistics are for year t_4 . See Handel and Kolstad (2015) for more information on the population, their information about insurance options, and the link between costs, information, surplus, and insurance choices. Note that we cannot provide the exact sample size for all employees and dependents at the firm, to preserve the anonymity of the firm (though we can discuss sample size of the specific samples used in our analysis).

Health Plan Characteristics
Family Tier

	PPO	HDHP*
Premium	\$0	\$0
Health Savings Account (HSA)	No	Yes
HSA Subsidy	-	[\$3,000-\$4,000]**
Max. HSA Contribution	-	\$6,250***
Deductible	\$0****	[\$3,000-\$4,000]**
Coinsurance (IN)	0%	10%
Coinsurance (OUT)	20%	30%
Out-of-Pocket Max.	\$0****	[\$6,000-\$7,000]**

* We don't provide exact HDHP characteristics to help preserve firm anonymity.

**Values for family coverage tier (2+ dependents). Single employees (or w/ one dependent) have $.4 \times (.8 \times)$ the values given here.

***Single employee legal maximum contribution is \$3,100. Employees over 55 can contribute an extra \$1,000 in 'catch-up.'

****For out-of-network spending, PPO has a very low deductible and out-of-pocket max. both less than \$400 per person.

Table F2: This table presents key characteristics of the two primary plans offered over time 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.




Figure F2: This figure presents the smoothed distribution of expected costs from the cost model for families in the primary sample. These costs are fully covered by the PPO and only partially by the HDHP.

in Section 3. See Handel and Kolstad (2015b) for a wider range of related specifications and estimates.

Primary Model		
Estimates		
	Model Estimate	95% CI
Average μ_γ	$8.6 \cdot 10^{-5}$	$[8.19 \cdot 10^{-5}, 2.23 \cdot 10^{-4}]$
Std. Dev. μ_γ	$1.4 \cdot 10^{-5}$	$[9.41 \cdot 10^{-6}, 4.41 \cdot 10^{-5}]$
Gamble Interp. of Average μ_γ	920.47	[822.51,924.23]
σ_γ	$2.2 \cdot 10^{-9}$	$[5.98 \cdot 10^{-6}, 1.55 \cdot 10^{-4}]$
σ_ϵ , HDHP	0.11	[1.58,666.04]
Benefits knowledge:		
Any incorrect	98.04	[-614.70,377.52]
Any 'not sure'	-467.48	[-1670.66,127.94]
Time cost hrs. X prefs:		
Time cost hrs.	-9.72	[-90.07,118.86]
... X Accept, concerned	-118.15	[-282.81,-55.79]
... X Dislike	-128.98	[-293.99,-70.02]
Provider networks:		
HSP network bigger	-594.38	[-1842.45,562.52]
PPO network bigger	-2362.85	[-3957.68,-1286.62]
Not sure	-201.81	[-937.44,303.21]
TME guess:		
Overestimate	62.98	[-810.72,704.28]
Underestimate	-208.30	[-1154.63,837.19]
Not sure	-688.91	[-1987.28,320.99]
Average Friction Effect	-1787.40	[-2148.63,-906.96]
σ Friction Effect	1303.64	[1264.29,2329.12]
Likelihood Ratio	379.54	
Test Stat for Frictions		

Table F3: This table presents our primary estimates of our empirical choice framework. The first column presents the actual point estimates while the second column presents the 95% CI derived from the bootstrapped standard errors. Here, positive friction values indicate greater willingness-to-pay for high-deductible care.

Simulations**Mistakes in Equilibrium**

	Low σ_c	Low σ_c	High σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	Low μ_s	Low μ_s	High μ_s
	Low σ_s	High σ_s	Low σ_s	High σ_s	High σ_s

Full Frictions

High μ_f , High σ_f	.23 (.18)	.31 (.27)	.47 (.45)	.36 (.33)	.11 (.07)
High μ_f , Low σ_f	.56 (.56)	.61 (.61)	.50 (.50)	.43 (.43)	.09 (.09)
Low μ_f , High σ_f	.39 (.13)	.39 (.15)	.35 (.24)	.35 (.20)	.35 (.06)
Low μ_f , Low σ_f	.26 (.13)	.21 (.11)	.02 (.01)	.06 (.04)	.17 (.04)

Half Frictions

High μ_f , High σ_f	.21 (.18)	.28 (.26)	.44 (.42)	.34 (.32)	.09 (.06)
High μ_f , Low σ_f	.35 (.35)	.30 (.30)	.50 (.50)	.42 (.42)	.09 (.09)
Low μ_f , High σ_f	.33 (.14)	.30 (.14)	.20 (.15)	.23 (.14)	.26 (.06)
Low μ_f , Low σ_f	.17 (.09)	.11 (.06)	.01 (0)	.02 (.01)	.09 (.03)

Table F4: This table presents the proportion of consumers making mistakes when purchasing coverage, given the equilibrium price, for a range of underlying population micro-foundations.

Simulations**Results for $\alpha = 0.5$**

	Low σ_c	Low σ_c	High σ_c	High σ_c	High σ_c
	Low μ_s	Low μ_s	Low μ_s	Low μ_s	High μ_s
	Low σ_s	High σ_s	Low σ_s	High σ_s	High σ_s

Half Frictions, % Purchase

High μ_f , High σ_f	0.94	0.92	0.86	0.9	0.97
High μ_f , Low σ_f	1	1	1	1	1
Low μ_f , High σ_f	0.63	0.57	0.34	0.41	0.72
Low μ_f , Low σ_f	0.61	0.44	0.06	0.11	0.86
No Frictions	0.45	0.39	0.06	0.11	0.91

Half Frictions % Surplus

High μ_f , High σ_f	0.95	0.95	0.87	0.94	0.97
High μ_f , Low σ_f	1	1	1	1	1
Low μ_f , High σ_f	0.68	0.66	0.42	0.56	0.76
Low μ_f , Low σ_f	0.71	0.60	0.10	0.25	0.94
No Frictions	0.56	0.55	0.10	0.23	0.95

Table F5: This table presents the proportion of consumers purchasing generous coverage (top half) and the proportion of first-best surplus achieved in the market (bottom half) for a range of underlying population micro-foundations and a partially effective friction-reducing policy ($\alpha = 0.5$).