

Estimating welfare in insurance markets using variation in prices*

Liran Einav, Amy Finkelstein, and Mark R. Cullen[†]

August 2009

Abstract. We provide a graphical illustration of how standard consumer and producer theory can be used to quantify the welfare loss associated with inefficient pricing in insurance markets with selection. We then show how this welfare loss can be estimated empirically using identifying variation in the price of insurance. Such variation, together with quantity data, allows us to estimate the demand for insurance. The same variation, together with cost data, allows us to estimate how insurer's costs vary as market participants endogenously respond to price. The slope of this estimated cost curve provides a direct test for both the existence and nature of selection, and the combination of demand and cost curves can be used to estimate welfare. We illustrate our approach by applying it to data on employer-provided health insurance from one specific company. We detect adverse selection but estimate that the quantitative welfare implications associated with inefficient pricing in our particular application are small, in both absolute and relative terms.

JEL classification numbers: C13, C51, D14, D60, D82, I11.

Keywords: Asymmetric information; adverse selection; health insurance; efficiency cost.

*We are grateful to Felicia Bayer, Brenda Barlek, Chance Cassidy, Fran Filpovits, Frank Patrick, and Mike Williams for innumerable conversations explaining the institutional environment of Alcoa, to Colleen Barry, Susan Busch, Linda Cantley, Deron Galusha, James Hill, Sally Vegso, and especially Marty Slade for providing and explaining the data, to Tatyana Deryugina, Sean Klein, Dan Sacks, and James Wang for outstanding research assistance, and to Larry Katz (the Editor), three anonymous referees, Kate Bundorf, Raj Chetty, Peter Diamond, Hanming Fang, David Laibson, Jonathan Levin, Erzo Luttmer, Jim Poterba, Dan Silverman, Jonathan Skinner, and numerous seminar participants for helpful comments. The data were provided as part of an ongoing service and research agreement between Alcoa, Inc. and Stanford, under which Stanford faculty, in collaboration with faculty and staff at Yale University, perform jointly agreed-upon ongoing and ad-hoc research projects on workers' health, injury, disability and health care, and Mark Cullen serves as Senior Medical Advisor for Alcoa, Inc. We gratefully acknowledge support from the NIA (R01 AG032449), the National Science Foundation grant #SES-0643037 (Einav), the Alfred P. Sloan Foundation (Finkelstein), and the John D. and Catherine T. MacArthur Foundation Network on Socioeconomic Status and Health, and Alcoa, Inc. (Cullen).

[†]Einav: Department of Economics, Stanford University, and NBER, leinav@stanford.edu; Finkelstein: Department of Economics, MIT, and NBER, afink@mit.edu; Cullen: Department of Internal Medicine, School of Medicine, Stanford University, mrcullen@stanford.edu.

1 Introduction

The welfare loss from selection in private insurance markets is a classic result in economic theory. It provides, among other things, the textbook economic rationale for the near-ubiquitous government intervention in insurance markets. Yet there has been relatively little empirical work devoted to quantifying the inefficiency that selection causes in a particular insurance market, or the welfare consequences of potential policy interventions in that market. This presumably reflects not a lack of interest in this important topic, but rather the considerable challenges posed by empirical welfare analysis in markets with hidden information.

Recently, there have been several attempts to estimate the welfare costs of private information in particular insurance markets, specifically annuities (Einav, Finkelstein, and Schrimpf 2009) and health insurance (Carlin and Town 2007; Bundorf, Levin, and Mahoney 2008; Lustig 2008). These papers specify and estimate a structural model of insurance demand that is derived from the choices of optimizing agents, and recover the underlying (privately known) information about risk and preferences. This allows for rich, out of sample, counterfactual welfare analysis. However, it requires the researcher to make critical assumptions about the nature of both the utility function and individuals' private information. These modeling choices can have non-trivial effects on the welfare estimates. Moreover, they are often specific to the particular market studied, making it difficult to meaningfully compare welfare estimates across markets or to readily adapt these approaches from one context to another.

Our objective in this paper is therefore to propose a complementary approach to empirical welfare analysis in insurance markets. We make fewer assumptions about the underlying primitives, yet impose enough structure to allow for meaningful welfare analysis. These fewer assumptions come at the cost of limiting our welfare analyses to only those associated with the pricing of existing contracts.

We start in Section 2 by showing how standard consumer and producer theory – familiar to any student of intermediate micro – can be applied to welfare analysis of insurance markets with selection. As emphasized by Akerlof (1970) and Stiglitz (1987) among others, the key feature of markets with selection is that firms' costs depend on which consumers purchase their products. As a result, insurer's costs are endogenous to price. Welfare analysis therefore requires not only knowledge of how demand varies with price, but also information on how changes in price affect the costs of insuring the (endogenous) market participants. We use these insights to provide a particular graphical representation of the welfare cost of inefficient pricing arising from selection. We view these graphs as providing helpful intuition, and therefore as an important contribution of the paper. The graphs illustrate, among other things, how the qualitative nature of the inefficiency depends on whether the selection is adverse or advantageous.

Our graphical analysis also suggests a straightforward empirical approach to welfare analysis of pricing in insurance markets. Section 3 shows how our framework translates naturally into a series of estimating equations, and discusses the data requirements. The key observation is that the same pricing variation that is needed to estimate the demand curve (or willingness to pay) in any

welfare analysis – be it the consequences of tax policy, the introduction of new goods, or selection in insurance markets – can also be used to estimate the cost curve in selection markets, i.e. how costs vary as the set of market participants endogenously changes. The slope of the estimated cost curve provides a direct test of the existence and nature of selection that – unlike the widely used “bivariate probit” test for asymmetric information (Chiappori and Salanie 2000) – is not affected by the existence (or lack thereof) of moral hazard. Specifically, rejection of the null hypothesis of a constant (i.e. horizontal) marginal cost curve allows us to reject the null hypothesis of no selection, while the sign of the slope of the marginal cost curve tells us whether the resultant selection is adverse (if marginal cost is increasing in price) or advantageous (if marginal cost is decreasing in price).

Most importantly, with both the demand and cost curves in hand, welfare analysis of inefficient pricing caused by any detected selection is simple and familiar. In the same vein, the estimates lend themselves naturally to welfare analysis of a range of counterfactual public policies that change the price of existing contracts. These include insurance mandates, subsidies or taxes for private insurance, and regulation of the prices that private insurers can charge.

Our approach has several attractive features. First, it does not require the researcher to make (often difficult to test) assumptions about consumer preferences or the nature of their ex ante information. As long as we accept revealed preference, the demand and cost curves are sufficient statistics for welfare analysis of the pricing of existing contracts. In this sense, our approach is similar in spirit to Chetty (2008) and Chetty and Saez (2008) who show how key ex-post behavioral elasticities are sufficient statistics for welfare analysis of the optimal level of public insurance benefits (see also Chetty (2009) for a more general discussion of the use of sufficient statistics for welfare analysis).

Second, our approach is relatively straightforward to implement, and therefore potentially widely applicable. In particular, while cost data are often quite difficult to obtain in many product markets (so that direct estimation of the cost curve is often a challenge), direct data on costs tend to be more readily available in insurance markets since they require information on accident occurrences or insurance claims, rather than insight into the underlying production function of the firm. In addition, the omnipresent regulation of insurance markets offers many potential sources for pricing variation needed to estimate the demand and cost curves. Third, the approach is fairly general as it does not rely on specific institutional details; as a result, estimates of the welfare cost of adverse selection in different contexts may be more comparable.

These attractive features are not without cost. As mentioned already, the chief limitation of our approach is that our analysis of the welfare cost of adverse selection is limited to the cost associated with inefficient pricing of a fixed (and observed) set of contracts. Our approach therefore does not allow us to capture the welfare loss that adverse selection may create by distorting the set of contracts offered, which in many settings could be large.¹ In the end of Section 3 we discuss in

¹A related limitation is that our approach forces us to rely on uncompensated (Marshallian) demand for welfare analysis. To account for income effects, we would either need to assume them away (by assuming constant absolute risk aversion) or to impose more structure and specify a full model of primitives that underlies the demand function.

some detail the settings where this limitation may be less prohibitive.

Analysis of the welfare effects of distortions in the contract space due to selection – or of counterfactual public policies that introduce new contracts – requires modeling and estimating the structural primitives underlying the demand and cost curves, and it is in this sense that we view our approach as complementary to a full model of these primitives. We note, however, that although such richer counterfactuals are feasible with a more complete model of the primitives, in practice the existing papers (mentioned above) that fully modeled these primitives have primarily confined themselves to welfare analyses of the pricing of existing contracts, as we do in this paper. This presumably reflects both researchers’ (understandable) caution in taking their estimates too far out of sample, as well as the considerable empirical and theoretical challenges to modeling the endogenous contract response (Einav, Finkelstein, and Levin, 2010). Perhaps similar reasons may also explain why many (although not all) government interventions in insurance markets tend to focus on the pricing of contracts through taxes and subsidies, regulations, or mandates.

The last part of the paper (Section 4) provides an illustration of our approach by applying it to the market for employer-provided health insurance in the United States, a market of substantial interest in its own right. The existing empirical evidence on this market is consistent with asymmetric information (see Cutler and Zeckhauser (2000) for a review). However, until recently there has been relatively little empirical work on the welfare consequences of the detected market failure. Cutler and Reber (1998) is a notable exception. Like us, they analyze selection in employer-provided health insurance, and, like us, they also estimate the demand curve. A key distinction, however, is that while they provide important and novel evidence of the existence of adverse selection in the market, they do not estimate the cost curve, which is crucial for welfare analysis.

We utilize rich individual-level data from Alcoa, Inc., a large multinational producer of aluminum and related products. We observe the health insurance options, choices, and medical insurance claims of its employees in the United States. We use the fact that, due to Alcoa’s organizational structure, employees doing similar jobs in different sections of the company are faced with different prices associated with otherwise identical sets of coverage options. We verify that pricing appears orthogonal to the characteristics of the employees that the managers setting these prices can likely observe. Using this price variation, we estimate that marginal cost is increasing in price, and thus detect adverse selection in this market. However, we estimate the welfare costs associated with the inefficient pricing created by adverse selection to be small. Specifically, we estimate that in a competitive market the annual efficiency cost of this selection would be just below \$10 per employee, or about 3% of the total surplus at stake from efficient pricing. By way of comparison, this estimated welfare cost is an order of magnitude lower than our estimate of the dead weight loss that would arise from monopolistic pricing in this market. We also estimate that the social cost of public funds for the price subsidy that would be required to move from the (inefficient) competitive equilibrium to the efficient outcome is about five times higher than our estimate of the welfare gain from achieving the efficient allocation. These results are robust across a range of alternative specifications.

It is extremely important to emphasize that there is no general lesson in our empirical findings.

Our estimates are specific to our population and to the particular health insurance choices they face. Nonetheless, at a conceptual level, our findings highlight the importance of moving beyond detection of market failures to quantifying their welfare implications. Our particular findings provide an example of how it is possible for adverse selection to exist, and to impair market efficiency, without being easily remediable through standard public policies.

2 Theoretical framework

2.1 Model

Setup and notation. We consider a situation in which a given population of individuals is allowed to choose from exactly two available insurance contracts, one that offers high coverage (contract H) and one that offers less coverage (contract L). As we discuss in more detail below, it is conceptually straightforward to extend the analysis to more than two contracts, but substantially complicates the graphical presentation. To further simplify the exposition, we assume that contract L is no insurance and is available for free, and that contract H is full insurance. These are merely normalizations and straightforward to relax; indeed we do so in our empirical application.

A more important assumption is that we take the characteristics of the insurance contracts as given, although we allow the price of insurance to be determined endogenously. As we discuss in more detail in Section 3 below, this seems a reasonable characterization of many insurance markets; it is often the case that the same set of contracts are offered to observably different individuals, with variation across individuals only in the pricing of the contracts, and not in offered coverage. Our analysis is therefore in the spirit of Akerlof (1970) rather than Rothschild and Stiglitz (1976), who endogenize the level of coverage as well.

We define the population by a distribution $G(\zeta)$, where ζ is a vector of consumer characteristics. A key aspect of the analysis is that we do not specify the nature of ζ ; it could describe multi-dimensional risk factors, consumers' ex-ante risk perception, and/or preferences. We denote the (relative) price of contract H by p , and denote by $v^H(\zeta_i, p)$ and $v^L(\zeta_i)$ consumer i 's (with characteristics ζ_i) utility from buying contracts H and L , respectively. Although not essential, it is natural to assume that $v^H(\zeta_i, p)$ is strictly decreasing in p and that $v^H(\zeta_i, p = 0) > v^L(\zeta_i)$. Finally, we denote the expected monetary cost associated with the insurable risk for individual i by $c(\zeta_i)$. For ease of exposition, we assume that these costs do not depend on the contract chosen, i.e. that there is no moral hazard. We relax this assumption in Section 2.4, where we show that allowing for moral hazard does not substantively affect the basic analysis.

Demand for insurance. We assume that each individual makes a discrete choice of whether to buy insurance or not. Since we take as given that there are only two available contracts and their associated coverages, demand is only a function of the (relative) price p . We assume that firms cannot offer different prices to different individuals. To the extent that firms can make prices depend on observed characteristics, one should think of our analysis as applied to a set of individuals that only vary in unobserved (or unpriced) characteristics. We assume that if individuals choose to buy

insurance they buy it at the lowest price at which it is available, so it is sufficient to characterize demand for insurance as a function of the lowest price p .

Given the above assumptions, individual i chooses to buy insurance if and only if $v^H(\zeta_i, p) \geq v^L(\zeta_i)$. We can define $\pi(\zeta_i) \equiv \max \{p : v^H(\zeta_i, p) \geq v^L(\zeta_i)\}$, which is the highest price at which individual i is willing to buy insurance. Aggregate demand for insurance is therefore given by

$$D(p) = \int 1(\pi(\zeta) \geq p) dG(\zeta) = \Pr(\pi(\zeta_i) \geq p), \quad (1)$$

and we assume that the underlying primitives imply that $D(p)$ is strictly decreasing, continuous, and differentiable.

Supply and equilibrium. We consider $N \geq 2$ identical risk neutral insurance providers, who set prices in a Nash Equilibrium (a-la Bertrand). Although various forms of imperfect competition may characterize many insurance markets, we choose to focus on the case of perfect competition as it represents a natural benchmark for welfare analysis of the efficiency cost of selection; under perfect competition, symmetric information leads to efficient outcomes, so that any inefficiency can be attributed to selection and does not depend on the details of the pricing model. We note however that it is straightforward to replicate the theoretical and empirical analysis for any other *given* model of the insurance market, including models of imperfect competition.

We further assume that when multiple firms set the same price, individuals who decide to purchase insurance at this price choose a firm randomly. We also assume that the only costs of providing contract H to individual i are the insurable costs $c(\zeta_i)$.² The foregoing assumptions imply that the average (expected) cost curve in the market is given by

$$AC(p) = \frac{1}{D(p)} \int c(\zeta) 1(\pi(\zeta) \geq p) dG(\zeta) = E(c(\zeta) | \pi(\zeta) \geq p). \quad (2)$$

Note that the average cost curve is determined by the costs of the sample of individuals who endogenously choose contract H . The marginal (expected) cost curve³ in the market is given by

$$MC(p) = E(c(\zeta) | \pi(\zeta) = p). \quad (3)$$

In order to straightforwardly characterize equilibrium, we make two further simplifying assumptions. First, we assume that there exists a price \bar{p} such that $D(\bar{p}) > 0$ and $MC(p) < p$ for every $p > \bar{p}$. In words, we assume that it is profitable (and efficient, as we will see soon) to provide insurance to those with the highest willingness to pay for it.⁴ Second, we assume that if there exists \underline{p} such that $MC(\underline{p}) > \underline{p}$ then $MC(p) > p$ for all $p < \underline{p}$. That is, we assume that $MC(p)$ crosses the

²Note that $c(\zeta_i)$ reflects only direct insurer claims (i.e. payout) costs, and not other administrative (production) costs of the insurance company. We discuss in Section 3.2 below how such additional costs can be incorporated into the analysis.

³Note that there could be multiple marginal consumers. Because price is the only way to screen in our setup, all these consumers will together average (point-by-point) to form the marginal cost curve.

⁴This assumption seems to hold in our application. Bundorf, Levin, and Mahoney (2008) make the interesting observation that there are contexts where it may not hold.

demand curve at most once.⁵ It is easy to verify that these assumptions guarantee the existence and uniqueness of equilibrium. In particular, the equilibrium is characterized by the lowest break-even price, that is:

$$p^* = \min \{p : p = AC(p)\}. \quad (4)$$

2.2 Measuring welfare

We measure consumer surplus by the certainty equivalent. The certainty equivalent of an uncertain outcome is the amount that would make an individual indifferent between obtaining this amount for sure and obtaining the uncertain outcome. An outcome with a higher certainty equivalent therefore provides higher utility to the individual. This welfare measure is attractive as it can be measured in monetary units. Total surplus in the market is the sum of certainty equivalents for consumers and profits of firms. Throughout we ignore any income effects associated with price changes.⁶

Denote by $e^H(\zeta_i)$ and $e^L(\zeta_i)$ the certainty equivalent of consumer i from an allocation of contract H and L , respectively; under the assumption that all individuals are risk averse, the willingness to pay for insurance is given by $\pi(\zeta_i) = e^H(\zeta_i) - e^L(\zeta_i) > 0$. We can write consumer welfare as

$$CS = \int [(e^H(\zeta) - p) 1(\pi(\zeta) \geq p) + e^L(\zeta) 1(\pi(\zeta) < p)] dG(\zeta) \quad (5)$$

and producer welfare as

$$PS = \int (p - c(\zeta)) 1(\pi(\zeta) \geq p) dG(\zeta). \quad (6)$$

Total welfare will then be given by

$$TS = CS + PS = \int [(e^H(\zeta) - c(\zeta)) 1(\pi(\zeta) \geq p) + e^L(\zeta) 1(\pi(\zeta) < p)] dG(\zeta). \quad (7)$$

It is now easy to see that it is socially efficient for individual i to purchase insurance if and only if

$$\pi(\zeta_i) \geq c(\zeta_i). \quad (8)$$

In other words, in a first best allocation individual i purchases insurance if and only if his willingness to pay is at least as great as the expected social cost of providing to him the insurance.⁷

⁵In the most basic economic framework of insurance the difference between $\pi(\zeta)$ and $MC(\zeta)$ is the risk premium, and is positive for risk averse individuals. If all individual are risk averse, $MC(\zeta)$ will never cross the demand curve. In practice, however, there are many reasons for such crossing. Those include, among others, loading factors on insurance, moral hazard, and horizontal product differentiation. As a result it may not be socially efficient for all individuals to have insurance, even if they are all risk averse.

⁶In a textbook expected utility framework, this is equivalent to assuming that the utility function exhibits constant absolute risk aversion (CARA). When the premium changes are small relative to the individual's income (as in the choice we study in our empirical application below) it seems natural to view CARA as a reasonable approximation. An alternative would be to fully specify the underlying utility function, from which income effects can be derived. This is one additional limitation of our simpler approach.

⁷Implicit in this discussion is that insurer claims $c(\zeta_i)$ represent the full social cost associated with allocating insurance to individual i . To the extent that this is not the case, for example due to positive or negative externalities associated with insurance or imperfections in the production of the underlying good that is being insured, our measure of welfare would have to be adjusted accordingly.

In many contexts (including our application below), price is the only instrument available to affect the insurance allocation. In such cases, achieving the first best may not be feasible if there are multiple individuals with different $c(\zeta_i)$'s who all have the same willingness to pay for contract H (see footnote 3). It is therefore useful to define a constrained efficient allocation as the one that maximizes social welfare subject to the constraint that price is the only instrument available for screening. Using our notation, this implies that it is (constrained) efficient for individual i to purchase contract H if and only if

$$\pi(\zeta_i) \geq E(c(\tilde{\zeta})|\pi(\tilde{\zeta}) = \pi(\zeta_i)). \quad (9)$$

That is, if and only if $\pi(\zeta_i)$ is at least as great as the expected social cost of allocating contract H to *all* individuals with willingness to pay $\pi(\zeta_i)$. We use this constrained efficient benchmark throughout the paper, and hereafter refer to it simply as the efficient allocation.⁸

2.3 Graphical representation

We use the framework sketched above to provide a graphical representation of adverse and advantageous selection. Although the primary purpose of doing so is to motivate and explain the empirical estimation strategy, an important ancillary benefit of these graphs is that they provide what we believe to be helpful intuition for the efficiency costs of different types of selection in insurance markets.

Adverse selection. Figure 1 provides a graphical analysis of adverse selection. The relative price (or cost) of contract H is on the vertical axis. Quantity (i.e. share of individuals in the market with contract H) is on the horizontal axis; the maximum possible quantity is denoted by Q_{max} . The demand curve denotes the relative demand for contract H . Likewise, the average cost (AC) curve and marginal cost (MC) curve denote the average and marginal *incremental* costs to the insurer from coverage with contract H relative to contract L .

The key feature of adverse selection is that the individuals who have the highest willingness to pay for insurance are those who, on average, have the highest expected costs. This is represented in Figure 1 by drawing a downward sloping MC curve. That is, marginal cost is increasing in price and decreasing in quantity. As the price falls, the marginal individuals who select contract H have lower expected cost than infra-marginal individuals, leading to lower average costs. The essence of the private information problem is that firms cannot charge individuals based on their (privately known) marginal cost, but are instead restricted to charging a uniform price, which in equilibrium implies average cost pricing. Since average costs are always higher than marginal costs, adverse selection creates under-insurance, a familiar result first pointed out by Akerlof (1970). This under-insurance is illustrated in Figure 1. The equilibrium share of individuals who buy contract H is

⁸See Greenwald and Stiglitz (1986) who analyze efficiency in an environment with a similar constraint. See also Bundorf, Levin, and Mahoney (2008) who investigate the efficiency consequences of relaxing this constraint. In a symmetric information case, the first best could be achieved by letting prices fully depend on $\pi(\zeta_i)$ and $c(\zeta_i)$.

Q_{eqm} (where the AC curve intersects the demand curve), while the efficient number is $Q_{eff} > Q_{eqm}$ (where the MC curve intersects the demand curve).

The welfare loss due to adverse selection is represented by the shaded region CDE in Figure 1. This represents the lost consumer surplus from individuals who are not insured in equilibrium (because their willingness to pay is less than the average cost of the insured population) but whom it would be efficient to insure (because their willingness to pay exceeds their marginal cost). One could similarly evaluate and compare welfare under other possible allocations. For example, mandating that everyone buy contract H generates welfare equal to the area ABE minus the area EGH. This can be compared to welfare at the competitive equilibrium (area ABCD), welfare at the efficient allocation (area ABE), welfare from mandating everyone to buy contract L (normalized to zero), or the welfare effect of policies that subsidize (or tax) the equilibrium price. The relative welfare rankings of these alternatives is an open empirical question. A primary purpose of the proposed framework is to develop an empirical approach to assessing welfare under alternative policy interventions (including the no intervention option).

Advantageous selection. The original theory of selection in insurance markets emphasized the possibility of adverse selection, and the resultant efficiency loss from under-insurance (Akerlof 1970; Rothschild and Stiglitz 1976). Consistent with this theory, the empirical evidence points to several insurance markets, including health insurance and annuities, in which the insured have higher average costs than the uninsured. However, a growing body of empirical evidence suggests that in many other insurance markets, including life insurance and long-term care insurance, there exists “advantageous selection.” Those with more insurance have lower average costs than those with less or no insurance. Cutler, Finkelstein, and McGarry (2008) provide a review of the evidence of adverse and advantageous selection in different insurance markets.

Our framework makes it easy to describe the nature and consequences of advantageous selection. Figure 2 provides a graphical representation. In contrast to adverse selection, with advantageous selection the individuals who value insurance the most are those who have, on average, the least expected costs. This translates to upward sloping MC and AC curves. Once again, the source of market inefficiency is that consumers vary in their marginal cost, but firms are restricted to uniform pricing, and in equilibrium price is based on average cost. However, with advantageous selection the resultant market failure is one of over-insurance rather than under-insurance (i.e. $Q_{eff} < Q_{eqm}$ in Figure 2), as was pointed out by de Meza and Webb (2001), among others. Intuitively, insurance providers have an additional incentive to reduce price, as the infra-marginal customers whom they acquire as a result are relatively good risks. The resultant welfare loss is given by the shaded area CDE, and represents the excess of MC over willingness to pay for individuals whose willingness to pay exceeds the average costs of the insured population. Once again, we can also easily evaluate welfare of different situations in Figure 2, including mandating contract H (the area ABE minus the area EGH), mandating contract L (normalized to zero), competitive equilibrium (ABE minus CDE), and efficient allocation (ABE).

Sufficient statistics for welfare analysis. These graphical analyses illustrate that the demand

and cost curves are sufficient statistics for welfare analysis of equilibrium and non-equilibrium pricing of existing contracts. In other words, different underlying primitives (i.e. preferences and private information as summarized by ζ) have the same welfare implications if they generate the same demand and cost curves.⁹

This in turn is the essence of our empirical approach. We estimate the demand and cost curves, but remain agnostic about the underlying primitives that give rise to them. As long as individuals' revealed choices can be used for welfare analysis, the precise *source* of selection is not germane for analyzing the efficiency consequences of the resultant selection, or the welfare consequences of public policies that change the equilibrium price.

The key to any counterfactual analysis that uses the approach we propose is that insurance contracts are taken as given, and only their prices vary. Thus, for example, the estimates generated by our approach can be used to analyze the effect of a wide variety of standard government interventions in insurance markets which change the price of insurance. These include mandatory insurance coverage, taxes and subsidies for insurance, regulations that outlaw some of the existing contracts, regulation of the allowable price level, or regulation of allowable pricing differences across observably different individuals. However, more structure and assumptions would be required if we were to analyze the welfare effects of introducing insurance contracts not observed in the data.

2.4 Incorporating moral hazard

Thus far we have not explicitly discussed any potential moral hazard effect of insurance. This is because moral hazard does not fundamentally change the analysis, but only complicates the presentation. We illustrate this by first discussing the baseline case in which we define contract H to be full coverage and contract L to be no coverage. Here, moral hazard has no effect on the welfare analysis. We then discuss the slight modification needed when we allow contract L to include some partial coverage.

With moral hazard, the expected insurable cost for individual i is now a function of his contract choice because coverage may affect behavior. We therefore define two (rather than one) expected monetary costs for individual i . We denote by $c^H(\zeta_i)$ individual i 's expected insurable costs under contract H relative to contract L , when he behaves as if he is covered by contract H . Similarly, we define $c^L(\zeta_i)$ to be individual i 's expected insurable costs under contract H relative to contract L , when he behaves as if he is covered by contract L . That is, $c^j(\zeta_i)$ always measures the incremental insurable costs under contract H compared to contract L , while the superscript j denotes the underlying behavior, which depends on coverage. We assume throughout that $c^H(\zeta_i) \geq c^L(\zeta_i)$; this inequality will be strict if and only if moral hazard exists. As a result, we now have two marginal cost curves, MC^H and MC^L , and two corresponding average cost curves, AC^H and AC^L (with MC^H and AC^H always higher than MC^L and AC^L , respectively).

⁹Note that we have placed no restrictions in Figures 1 or 2 on the nature of the underlying consumer primitives ζ_i . Individuals may well differ on many unobserved dimensions concerning their information and preferences. Nor have we placed any restriction on the nature of the correlation across these primitives.

In contrast to the selection case, a social planner generally has no potential comparative advantage over the private sector in ameliorating moral hazard (i.e. in encouraging individuals to choose socially optimal behavior). Our welfare analysis of selection therefore takes any moral hazard effect as given. We investigate the welfare cost of the inefficient pricing associated with selection or the welfare consequences of particular public policy interventions given any existing moral hazard effects, just as we take as given other features of the environment that may affect willingness to pay or costs.

In order to explicitly recognize moral hazard in our foregoing equilibrium and welfare analysis one can simply replace $c(\zeta_i)$ everywhere above with $c^H(\zeta_i)$, and obtain the same results. Recall, as emphasized earlier, that the cost curve is defined based on the costs of individuals who endogenously buy contract H (see equation (2)); in the new notation their costs are given by $c^H(\zeta_i)$ since they are covered by contract H (and behave accordingly). Thus, $c^L(\zeta_i)$ is largely irrelevant. The intuition from the firm perspective is clear: the insurer's cost is only affected by the behavior of insured individuals, and not by what their behavior would be if they were not insured. From the consumer side $c^L(\zeta_i)$ does matter. However, it matters only because it is one of the components that affect the willingness to pay for insurance. As we showed already, willingness to pay (π) and cost to the insurer (c^H) are sufficient statistics for the equilibrium and welfare analysis. Both can be estimated without knowledge of $c^L(\zeta_i)$. Therefore, as long as moral hazard is taken as given, it is inconsequential to break down the willingness to pay for insurance to a part that arises from reduction in risk and a part that arises from a change in behavior.

The one substantive difference once we allow for moral hazard is that the assumption that contract L involves no coverage is no longer inconsequential. Once contract L involves some partial coverage, it is no longer the case that all potential moral hazard effects of contract H on insurable expenditures are internalized by the provider of contract H through their impact on c^H . To see this, we first note that when contract L involves some coverage, the market equilibrium could be thought of as one in which firms offering contract H only compete on the incremental coverage in excess of L .¹⁰ Welfare analysis of the allocation of contract H must now account for the potential negative externality that coverage by contract H inflicts on the insurer providing contract L (through increased cost). This conceptual point does not pose practical difficulties for our framework. With estimates of the moral hazard effect, the welfare gain of providing contract H to individual i is simply smaller by the amount of the increased insurable costs for the provider of contract L that are associated with the change of behavior. As we discuss in more detail in Section 3, our approach points to a natural way by which moral hazard can be estimated (and therefore incorporated into the welfare analysis if needed, when contract L involves some partial coverage).

¹⁰One natural example is that of contract L as the public health insurance program Medicare and contract H as the supplemental private Medigap insurance that covers some of the costs not covered by Medicare.

3 Estimation

3.1 The basic framework

Applying our framework to estimating welfare in an insurance market requires data that allows estimation of the demand curve $D(p)$ and the average cost curve $AC(p)$. The marginal cost curve can be directly backed out from these two curves and does not require further estimation. To see this, note that

$$MC(p) = \frac{\partial TC(p)}{\partial D(p)} = \frac{\partial (AC(p) \cdot D(p))}{\partial D(p)} = \left(\frac{\partial D(p)}{\partial p} \right)^{-1} \frac{\partial (AC(p) \cdot D(p))}{\partial p}. \quad (10)$$

With these three curves – $D(p)$, $AC(p)$, and $MC(p)$ – in hand, we can straightforwardly compute welfare under various allocations, as illustrated in Figures 1 and 2.

As is standard, estimating the demand curve requires data on prices and quantities (i.e. coverage choices), as well as identifying price variation that can be used to trace out the demand curve. This price variation has to be exogenous to unobservable demand characteristics. To estimate the $AC(p)$ curve we need, in addition, data on the expected costs of those with contract H , such as data on subsequent risk realization and how it translates to insurer costs. With such data we can then use the same variation in prices to trace out the $AC(p)$ curve. Because expected cost is likely to affect demand, any price variation that is exogenous to demand is also exogenous to insurable cost. That is, we do not require a separate source of variation.

With sufficient price variation, no functional form assumptions are needed for the prices to trace out the demand and average cost curves. For example, if the main objective is to estimate the efficiency cost of inefficient pricing arising from selection, then price variation that spans the range between the market equilibrium price (point C in Figures 1 and 2) and the efficient price (point E) allows us to estimate the welfare cost of the inefficient pricing associated with selection (area CDE) without making any restrictions on the shape of the demand or average cost curves. With pricing variation that does not span these points, the area CDE can still be estimated, but will require some extrapolation based on functional form assumptions.

3.2 Extensions

As mentioned, the basic framework we described in Section 2 made a number of simplifying assumptions for expositional purposes which do not limit the ability to apply this approach more broadly. It is straightforward to apply the approach to the case where contract H provides less than full coverage and/or where contract L provides some coverage. We discuss a specific example of this in our application below. In such settings we must simply be clear that the cost curve of interest is derived from the average *incremental* costs to the insurance company associated with providing contract H rather than providing contract L . For the welfare analysis, we must also be sure to incorporate any moral hazard effects of contract H on the costs to the insurers providing contract L . We discussed above conceptually how to adjust the welfare analysis; later in this section we describe how to estimate the moral hazard effect of contract H .

Likewise, while it was simpler to present the graphical analysis with only two coverage options, the approach naturally extends to more than two contracts. The data requirements would simply extend to having price, quantity, and costs for each contract, as well as pricing variation across all relevant relative prices, so that the entire demand and average cost systems can be estimated. Specifically, with N available contracts, one could normalize one of these contracts to be the reference contract, define incremental costs (and price) of each of the other contracts relative to the reference contract, and estimate a system $D(p)$ and $AC(p)$, where demand, prices, and average costs are now $N - 1$ dimensional vectors. As in the two-contract case, competitive equilibrium (defined by each contract breaking even) will be given by the vector of prices that solves $p = AC(p)$. From the estimated systems $D(p)$ and $AC(p)$ one can also back out the system of marginal costs $MC(p)$ which defines the marginal costs associated with each price vector. We can then solve $p = MC(p)$ for the efficient price vector and integrate $D(p) - MC(p)$ over the (multi-dimensional) difference between the competitive and the efficient price vectors to obtain the welfare cost of the inefficient pricing associated with selection.¹¹

Finally, we note that the estimated demand and cost curves are sufficient statistics for welfare analysis of equilibrium allocations of existing contracts generated by models other than the one we have sketched. This includes, for example, welfare analysis of other equilibria, such as those generated by imperfect competition rather than our benchmark of perfect competition. It also includes welfare analysis of markets with other production functions, which may include fixed or varying administrative costs of selling more coverage, rather than our benchmark of no additional costs beyond insurable claims. This is because, as the discussion of estimation hopefully makes clear, we do not use assumptions about the equilibrium or the production function to estimate the demand and cost curves. An assumption of a different equilibrium simply requires calculation of welfare relative to a different equilibrium point (point C in the graphs). Similarly, if one has external information (or beliefs) about the nature of the production function, one can use this to shift or rotate the estimated cost curve, and calculate the new equilibrium and efficient points.

3.3 A direct test of selection

Although the primary focus of our paper is on estimating the welfare cost of inefficient pricing associated with selection, our proposed approach also provides a direct test for the existence and nature of selection. This test is based on the slope of the estimated marginal cost curve. A rejection of the null hypothesis of a constant marginal cost curve allows us to reject the null of no selection.¹² Moreover, the sign of the slope of the estimated marginal cost curve informs us of

¹¹While conceptually straightforward, implementation of our approach with more than two contracts will likely encounter, in practice, a number of subtle issues. For example, with multiple contracts the systems $AC(p) = p$ or $MC(p) = p$ may have more scope for multiple or no solutions, and the definition of “adverse selection” or “advantageous selection” may now be more subtle (see Einav, Finkelstein, and Levin (2010) for more discussion of this latter point). In addition, from an empirical standpoint, estimating an entire demand and cost systems may be more challenging (e.g., in terms of the variation required) than estimating one-dimensional demand and cost curves.

¹²Using the terminology we defined in Section 2.2, a flat marginal cost curve implies that the equilibrium outcome is constrained efficient. It does not however imply that the equilibrium is first best. Finkelstein and McGarry (2006)

the nature of any selection; a downward sloping marginal cost curve (i.e. a cost curve declining in quantity and increasing in price) indicates adverse selection, while an upward sloping curve indicates advantageous selection. This is a useful test, since detecting the existence of selection is a necessary precursor to analysis of its welfare effects.

Importantly, our “cost curve” test of selection is unaffected by the existence (or lack thereof) of moral hazard. This is a distinct improvement over the influential “bivariate probit” (a.k.a. “positive correlation”) test of Chiappori and Salanie (2000) which has been widely used in the insurance literature. This test, which compares realized risks of individuals with more and less insurance coverage, jointly tests for the existence of either selection or moral hazard (but not for each separately). Identifying price variation – which is not required for the “positive correlation” test – is the key to our distinct test for selection. It allows us to analyze how the risk characteristics of the sample who selects a given insurance contract vary with the price of that contract.

To see why our cost curve test is not affected by any potential moral hazard, note that the AC curve is estimated using the sample of individuals who choose to buy contract H at a given price. As we vary price we vary this sample, but everyone in the sample always has the same coverage. Since by construction the coverage of individuals in the sample is fixed, our estimate of the slope of the cost curve (our test of selection) is not affected by moral hazard (which determines how costs are affected as coverage changes). Of course, part of the selection reflected in the slope of the cost curve may reflect selection based on differences across individuals in the anticipated impact of coverage on costs (i.e. the moral hazard effect of coverage). We still view this as a selection effect, representing selection into contracts based on the anticipated incentive effects of these contracts.

3.4 Estimating moral hazard

Our framework also allows us to test for and quantify moral hazard. One way to measure moral hazard is by the difference between $c^H(\zeta_i)$ – individual i ’s expected insurable cost when he is covered by contract H – and $c^L(\zeta_i)$ – individual i ’s expected insurable cost when he is covered by contract L . That is, $c^H(\zeta_i) - c^L(\zeta_i)$ is the moral hazard effect from the insurer’s perspective, or the increased cost to the insurer from providing contract H that is attributable to the change in behavior of covered individuals. We already discussed how identifying price variation can be used to estimate the AC and MC curves, which we denote by AC^H and MC^H when moral hazard is explicitly recognized. With data on the costs of the uninsured (or less insured, if contract L represents some partial coverage), we can repeat the same exercise to obtain an estimate for AC^L and MC^L . That is, we can use the same identifying price variation to estimate demand for contract L and to estimate the AC^L curve from the (endogenously selected) sample of individuals who choose contract L . We can then back out the MC^L curve analogously to the way we back out the MC^H curve, using of course the demand curve for contract L and the AC^L curve (rather than the demand for contract H and the AC^H curve) in translating average costs into marginal costs (see equation (10)). The

present evidence on an insurance market that may exhibit a flat cost curve (no selection) but does not achieve the first best allocation.

(point-by-point) vertical difference between MC^H and MC^L curves provides an estimate of moral hazard. A test of whether this difference is positive is a direct test for moral hazard, which is valid whether adverse selection is present or not.¹³

Of course, it is not a new observation that an exogenous shifter of insurance coverage (which in our context comes from pricing) facilitates the estimation of moral hazard. However, our proposed approach to estimating moral hazard (compared to, say, a more standard instrumental variable framework) allows us to estimate (with sufficiently rich price variation) heterogeneous moral hazard effect and to see how moral hazard varies across individuals with different willingness to pay $\pi(\zeta_i)$, or different expected costs $c^H(\zeta_i)$.

3.5 Applicability

In the next section we turn to a specific application of our proposed framework, which illustrates the mechanics of the approach as well as produces results that may be of interest in their own right. Here we discuss more generally the types of settings in which our approach might be applicable.

There are two main requirements that need to be met in order to sensibly use our approach. First, it has to be feasible to credibly estimate the demand and cost curves. This requires both data on insurance prices, quantities, and insurer's costs, as well as identifying variation in prices. The required data elements of insurance options and choices and subsequent risk realization are not particularly stringent; researchers have already demonstrated considerable success in a wide range of insurance markets in obtaining such data.¹⁴ Indeed, a nice feature of welfare analysis in insurance markets is that cost data are much easier to obtain than in many other markets, since they involve information on accident occurrences or insurance claims, rather than insight into the underlying production function of the firm.

Finding identifying variation in prices is a considerably stronger empirical hurdle, although the near-ubiquitous regulation of insurance markets provide numerous potential opportunities. Recall that while our application below assumes that prices are set exogenously to unobservable demand (and cost) characteristics, alternative research designs that isolate credible identifying variation would do. This could include an instrumental variable approach, a regression discontinuity design, and so forth. For example, state regulations of private insurance markets have created variation in the prices charged to different individuals at a point in time as well as over time (Blackmon and Zeckhauser 1991; Buchmueller and DiNardo 2002; Bundorf and Simon 2006). Tax policy is another useful potential source of pricing variation. For example, a large literature has documented

¹³The exercise we have just described would provide an estimate of the moral hazard effect from the insurer's perspective. One might be interested in other measures of moral hazard, such as the effect of insurance on total spending rather than insurer costs. The test of moral hazard can be applied in the same manner using other definitions of $c(\zeta_i)$. The same statement of course applies to our "cost curve" selection analysis; for the purpose of analyzing equilibrium and market efficiency, we have estimated selection from the insurer perspective, but again the approach could be used to measure selection on any other outcome of interest.

¹⁴Examples include auto and homeowner's insurance (Sydnor 2006; Cohen and Einav 2007; Barseghyan, Prince, and Teitelbaum 2008), annuities (Finkelstein and Poterba 2004), long term care insurance (Finkelstein and McGarry 2006), health insurance (Eichner, McClellan, and Wise 1998), and many others.

(and used) the substantial variation in the tax subsidy for employer-provided health insurance (see Gruber (2002) for a review). Beyond the opportunities provided by public policy, researchers have also found useful pricing variation stemming from field experiments (Karlan and Zinman 2009) and specific idiosyncrasies of firm pricing behavior.¹⁵ More generally, common instruments used in demand analysis, such as changes in the competitive environment (Lustig 2008) or perhaps shifters in the administrative costs of handling claims, could serve as the requisite source for identifying price variation. The validity of this variation for identification is of course a key issue, which can and should be evaluated in specific applications. Indeed, we see the transparency of our approach in this regard as an important attraction.

The second key requirement for applying our proposed framework stems from its focus on inefficient pricing. Given that it is designed to estimate the welfare consequences from pricing of *existing* contracts, it is best suited to settings in which the market or public policy response to asymmetric information will primarily manifest itself through pricing of observed contracts rather than other aspects of contract design. We note that a pricing response also covers the elimination of certain contracts or mandating a specific (observed) contract, which is of course equivalent to increasing a subset of the prices to infinity. What our approach cannot accommodate is a response that leads to the introduction of new contracts, which were not previously observed.

How closely a given setting fits this bill needs to be evaluated on a case by case. Perhaps the ideal setting is one in which regulation (or some other constraint) explicitly prevents firms from redesigning contracts. While rare, examples exist. One such case is the (limited) set of contracts that can be offered in the Medigap market, the private health insurance that supplements Medicare. Since 1992, these contracts have been set by national regulation: private firms may decide which of the specified contracts to offer and at what price, but they cannot design and introduce new contracts (see, e.g., Fox, Rice, and Alexih 1995). A related example is the application we discuss below in which company headquarters design the coverage options and print the brochures that describe them, while different subsidiaries are allowed (some) choice over the relative pricing of these options.

A likely more common setting that doesn't quite fit this ideal standard but may come sufficiently close is the practice in many markets to first settle on the contract design, and then adjust only prices over time and across individuals. For example, the Medicare Part D market (for subsidized prescription drug coverage for the elderly) divides the country into 34 geographical markets. Those providers that operated in multiple markets (and most of them do) have designed and advertised a single (national) set of coverage plans (in terms of formularies, deductible, cost sharing, etc.), and only adjusted their prices by region (Keating 2007) Similarly, in annuity markets companies offer identical sets of contracts (in terms of tilt of payments and guaranteed payment features), with only the annuity rates varying with the annuitant mortality profile (Finkelstein and Poterba 2002).

¹⁵Examples include firm experimentation with their pricing policy (Cohen and Einav 2007), discrete pricing policy changes (Adams, Einav, and Levin 2009), idiosyncratic pricing decisions made by human resource managers (Cutler and Reber 1998), and the non-linearities and discontinuities associated with rules firms use to risk adjust individuals' premiums (Abbring, Chiappori, and Pinquet 2003; Israel 2004).

It is our impression that life insurance and long-term care insurance contracts tend to operate in a similar way. Such behavior is consistent with the notion that adjusting prices is less costly than adjusting other contract features, possibly due to issues related to marketing and advertising, standardization, or norms. Whatever such mechanism may be, it lends itself to our framework. That is, in these settings the most likely response to non-drastic changes in information structure would likely be primarily manifested along the price dimension.

4 Empirical illustration: employer-provided health insurance

4.1 Data and environment

We illustrate the approach we have just outlined using individual-level data from 2004 on the U.S.-based employees (and their dependents) at Alcoa, Inc. The primary purpose of the application is to show how the theoretical framework can be mapped into empirical welfare estimates. We view the direct link between the theoretical framework and the empirical estimates – and the resulting transparency this provides for evaluating the strengths and weaknesses of the empirical results – as a key strength of our approach.

In 2004 Alcoa had approximately 45,000 active employees in the U.S. working at about 300 different job sites in 39 different states. At that time, in an effort to control healthcare spending, Alcoa introduced a new set of health insurance options to virtually all its salaried employees and about one-half of its hourly employees. We analyze the choices and claims of employees offered the new set of options in 2004.¹⁶

The data contain the menu of health insurance options available to each employee, the employee premium associated with each option, the employee’s coverage choice, and detailed claim-level information on all his (and any covered dependents’) medical expenditures during the coverage period.¹⁷ Crucially, as we discuss below, the data contain plausibly exogenous variation in the prices of the insurance contracts offered to otherwise similar employees within the company. Finally, the data contain rich demographic information, including the employee’s age, race, gender, annual earnings, job tenure at the company, and the number and ages of other insured family members. We suspect that we observe virtually everything about the employee that the administrators setting insurance premiums can observe without direct personal contact, as well as some characteristics that the price setters might not be able to observe (such as detailed medical expenditure information from previous years; this information is administered by a third party). This is important because it allows us to examine whether the variation in prices across employees appears correlated with the employee characteristics that could potentially influence the price setters’ decisions.

¹⁶Over the subsequent several years, most of the remaining hourly employees were transitioned to the new health insurance options as their union contracts expired. The variation over time in the contracts offered is not well suited to the approach developed here, which relies on variation in the pricing of the same set of contract offerings. Busch et al. (2006) study the effect of the change in plan options between 2003 and 2004 on the use of preventive care.

¹⁷Health insurance choices are made during the “open enrollment” period at the end of 2003 and apply for all of 2004. We also observe medical expenditure in 2003 if the employee worked at the company for all of 2003.

We restrict our baseline analysis to a subsample of employees for whom the pricing variation is cleaner and the setting follows more closely the theoretical framework. Our baseline sample consists of 3,779 salaried employees with family coverage who chose one of the two modal health insurance choices: a higher and a lower level of PPO coverage (we refer to these hereafter as contract H and contract L and provide more details about them in Section 4.3 below). The online appendix provides many more details about these sample restrictions, provides results for other coverage tiers, and addresses concerns of sample selection.

4.2 Variation in prices

Company structure as the source of variation. An essential element in the analysis is that there is variation across employees in the relative price they face for contract H , and that this variation is unrelated to the employees' willingness to pay for contract H or to his insurable costs. We believe that Alcoa's business structure provides a credible source of such pricing variation across different employees in the company.

In 2004, as part of the new benefit design, company headquarters offered a set of seven different possible pricing menus for employee benefits. The coverage options are the same across all the menus, but the prices (employee premiums) associated with these options vary. For our purposes, the key element of interest is the incremental (annual) premium the employee must pay for contract H relative to contract L , $p = p_H - p_L$. We refer to this incremental premium as the "price" in everything that follows. There were six different values of p in 2004 (as two of the seven menus were identical in this respect), ranging (for family coverage) from \$384 to \$659.¹⁸

Which price menu a given employee faces is determined by the president of his business unit. Alcoa is divided into approximately forty business units. Each business unit has essentially complete independence to run their business in the manner they see fit, provided that they do so ethically and safely, and at or above the company's normal rate of return. Failure on any of these dimensions can result in the replacement of the unit's president. Business units are typically organized by functionality – such as primary metals, flexible packaging materials, rigid packaging materials, or home exterior – and are independent of geography. There are often multiple business units in the same state. The number of active employees in a business unit ranges from the low teens (in "government affairs") to close to 6,000 (in "primary metals"). The median business unit has about 500 active employees. The business unit president may choose different price menus for employees within his unit based on their location (job site) and their employment type (salaried or hourly employee and, if hourly, which union if at all the employee is in).

As a result of this business structure, employees doing the same job in the same location may face different prices for their health insurance benefits due to their business unit affiliation. A priori, it

¹⁸The annual pre-tax employee premium for contract H was around \$1,500 for family coverage, although of course it ranged across the different menus. The incidence of being offered a menu with a lower average price level (across different options) may well be passed on to employees in the form of lower wages (Gruber 1994). This is one additional reason why it is preferable to focus the analysis on the difference in premiums for the different coverage options, rather than the level of premiums.

struck us as more plausible that the pricing variation across salaried employees in different business units is more likely to be useful for identification – reflecting idiosyncratic characteristics of the business unit presidents rather than differences in the demand or costs of salaried employees in the different business units – than the pricing variation across hourly employees. This is because many of the jobs that salaried employees do are quite similar across business units. Thus, for example, accountants, paralegals, administrative assistants, electrical engineers, or metallurgists working in the same state may face different prices because their benefits were chosen by the president of the “rigid packaging” business unit, rather than by the president of “primary metals.” By comparison, the nature of the hourly employees’ work (which often involves the operation of particular types of machinery) is more likely to differ across different units, and may depend on what the business unit is producing. For example, the work of the potroom operators stirring molten metal around in large vats in the “primary metals” business unit is likely to be different from the work of the furnace operators in the “rigid packaging” unit.

Examination of assumption of exogenous pricing. The available data appear consistent with this basic intuition. Table 1 compares mean demographic characteristics of employees in our baseline sample who face different prices. In general, the results look quite balanced. There is no substantive or statistically significant difference across employees who face different prices in average age, fraction male, fraction white, average (log) wages, average age of spouse, number of covered family members, age of the youngest child, or in the average (log) 2003 medical spending. The two possible exceptions to this general pattern are average job tenure and average (log) 2003 medical expenditures when restricted to employees in the most common plan in 2003 (to avoid potential differences in spending arising from moral hazard effects of different 2003 coverages).¹⁹ A joint F – test of all of the coefficients leaves us unable to reject the null that they are jointly uncorrelated with price.²⁰ Inference is similar when we include state fixed effects or extend the sample to include all coverage tiers (rather than family coverage only) or all salaried employees (rather than just the two-thirds who choose the two modal coverage options).

Ancillary support of the quantitative evidence we have just described comes from our qualitative investigation into benefit selection at Alcoa in 2004. Importantly, this was the first year ever that business unit presidents had the opportunity to make decisions regarding the relative prices of insurance contracts for their employees. Therefore, while one might suspect that over time their price selection might become more sophisticated with respect to demand or expected costs (which would invalidate our identification assumption), in the first year the decision makers had relatively little information or experience to go by. Relatedly, the new benefit system represented the first time in the company’s history that it was possible to charge employees a substantial incremental

¹⁹The p – value on each of these two variables is 0.08. We should note, of course, that when testing 10 different variables the p – value should be adjusted upward to take account of the multiple hypothesis testing, so that the p – values we report are too small.

²⁰When we examine the eight contemporaneous characteristics we obtain an F – stat of 1.71 (p – value = 0.14). When we also include 2003 spending for those in the same plan as a ninth covariate (so that our sample size falls by about 25 percent) we obtain an F – stat of 0.98 (p – value = 0.48).

price for greater health insurance coverage. Our discussions with the company suggested that many business unit presidents were (at least initially) philosophically opposed to charging employees much for (generous) health insurance coverage, which may explain why (as we show below in Table 2) about three quarters of the salaried employees ended up facing the lowest possible incremental price that the business unit presidents were allowed to choose. Perhaps because of this, after 2004 Alcoa headquarters no longer gave the business unit presidents a choice on benefit prices, and chose a (uniform) pricing structure with a higher price than any of the options available in 2004.

Interestingly, the story looks very different for hourly employees. A similar analysis of covariates for hourly employees suggests statistically significant differences across employees who face different prices. As noted, this is not surprising given the institutional environment, and motivates our sample restriction to salaried employees. Indeed, the fact that prices for hourly employees are not uncorrelated with employee characteristics is somewhat reassuring; in a large for-profit company, it makes sense to expect clear differences in employee characteristics to be reflected in the prices chosen. It may be that when there was more at stake (in terms of cost differences across employees) the business unit presidents paid more attention to setting prices and less to their idiosyncratic philosophical views. It is also possible – although we have no direct evidence for this – that the business unit presidents had fundamentally different objectives in setting prices for hourly and for salaried employees.

Thus, while we would of course prefer to be able to isolate the precise source of our pricing variation, we are nonetheless reassured by both the quantitative and qualitative evidence that the prices faced by salaried employees appear uncorrelated with their predictors of demand or costs. Of course, we are able to only examine whether prices are correlated with observable differences across salaried employees. We cannot rule out potential unobservable differences, for example in the “culture” of the business unit, which could potentially affect price setting and be correlated with either demand or costs.

4.3 Empirical strategy and relationship with the theoretical framework

As before, we denote by $p_i = p_i^H - p_i^L$ the relative price employee i faces, where p_i^j is employee i 's annual premium if she chose coverage j . We define D_i to be equal to 1 if employee i chooses contract H and 0 if employee i chooses contract L . Finally, we let m_i be a vector representing total medical expenditures of employee i and any covered family members in 2004.

Coverage characteristics and construction of the cost variable. In our theoretical discussion in Section 2 we defined (for simplicity) contract H to be full coverage and contract L to be no coverage. As a result we could refer to c_i as the *total* cost to the insurance company from covering employee i . When contract H is not a full coverage and contract L provides some partial coverage, the relevant cost variable (denoted c_i) is defined as the *incremental* cost to the insurer from providing contract H relative to providing contract L , holding m_i fixed. Specifically, let $c(m_i; H)$ and $c(m_i; L)$ denote the cost to the insurance company from medical expenditures m_i under contracts H and L , respectively. The incremental cost is then given by $c_i \equiv c(m_i) = c(m_i; H) - c(m_i; L)$. The *AC* curve is computed

by calculating the average c_i for all individuals who choose contract H at a given relative price p (see equation (2)) and estimating how this average c_i varies as the relative price varies. We can observe $c(m_i; H)$ directly in the data, but $c(m_i; L)$ must be computed counterfactually using the claims data and the coverage rules of contract L . For consistency, we calculate both $c(m_i; H)$ and $c(m_i; L)$ from plan rules.

Construction of c_i requires detailed knowledge of each plan’s benefits as well as individuals’ realized medical claims. This allows us to construct the cost to the insurance company of insuring medical expenditures m_i under any particular plan j . The two contracts we focus on vary only in their consumer cost-sharing rules. Specifically, contract L coverage has higher deductibles and higher out-of-pocket maximums.²¹ The data are quite detailed and the plan rules are fairly simple, allowing us to calculate $c(m_i; j)$ with a great deal of accuracy. For example, for individuals with contract H the correlation between their actual (observed) share of out-of-pocket spending (out of total expenditure) and our constructed share is over 0.97. The online appendix provides more detail on our calculation of c_i .

Figure 3 presents the major differences in consumer cost sharing between the two coverage options. Cost sharing rules differ depending on whether spending is in-network or out-of-network. Figure 3(a) shows the annual out-of-pocket spending (on the vertical axis) associated with a given level of total medical spending m (on the horizontal axis) for each coverage option, assuming the medical spending is in-network. In network, contract H has no deductible while contract L has a \$500 deductible. Both contracts have a 10% coinsurance rate, and the out-of-pocket maximum is \$5,000 for contract H and \$5,500 for contract L . Figure 3(b) presents the analogous graph for out-of-network spending, which has higher cost sharing requirements under both plans. Although the vast majority of spending (96%) occurs in network, about 25% of the individuals in our baseline sample file at least one claim out of network, making the out-of-network coverage an important part of the analysis.²²

Figures 3(c) and 3(d) show the implied difference in out-of-pocket spending between contracts H and L , for a given level of annual medical expenditure m_i . Figure 4 presents the empirical distribution of the constructed c_i variable. The distribution of c_i reflects the various kinks in the coverage plans presented in Figure 3. The most visible example is that about two thirds of the individuals in our baseline sample have $c_i = 450$. This represents individuals who had between \$500 and \$50,000 in-network (total) medical expenditures and less than \$500 out-of-network expenditures.²³

The nature of the plan differences is important for understanding the margin on which we may detect selection (or moral hazard). Empirically, because only few people spend anywhere

²¹The plans are similar in all other features, such as the network definition and the benefits covered. As a result, we do not have to worry about differences between contracts H and L in plan features that might differ in unobservable ways across employees (for example, differences in providers, the relative network quality, and so forth).

²²There is no interaction between the in-network and out-of-network coverages. Each deductible and out-of-pocket maximum must be satisfied separately.

²³Note also that, as emphasized by Figure 4, because our cost variable captures the incremental cost of increased coverage (rather than total medical expenditures) it is not heavily influenced by outliers (catastrophic expenditures). Indeed, as shown in Figure 3, plan rules essentially cap incremental costs at \$1,500.

close to the out-of-pocket maximum of either contract, the difference in insurer's cost between the plans is primarily attributable to differences in the deductible. In terms of selection, this suggests that the differences in the plans could matter for the insurance choice of anyone with positive expected expenditures, and is increasing as expected expenditures increase. In terms of moral hazard, this suggests that if individuals are forward looking and have perfect foresight then differences in behavior for people covered by the different plans should be limited to the small percentage (9%) of employees who have total medical expenditures that are less than the contract L deductible.

Baseline estimating equations. For our baseline specification, we estimate the demand and average cost functions using OLS, assuming that the demand and cost curves are each linear in prices. That is, we estimate the following two equations

$$D_i = \alpha + \beta p_i + \epsilon_i \quad (11)$$

$$c_i = \gamma + \delta p_i + u_i \quad (12)$$

where D_i , c_i , and p_i are defined earlier. In all regressions, we adjust the standard errors to allow for an arbitrary variance-covariance matrix within each state. This is to allow for potential correlation in the residuals of the demand or cost equations across salaried employees in the same state. Following the theoretical framework, the demand equation is estimated on the entire sample, while the (average) cost equation is estimated on the sample of individuals who (endogenously) choose contract H .

Using the point estimates from the above regressions, we can construct our predicted demand and average cost curves and other estimates of interest. Following equation (10), the marginal cost curve is given by

$$MC(p) = \frac{1}{\beta} \left(\frac{\partial(\alpha + \beta p)(\gamma + \delta p)}{\partial p} \right) = \frac{1}{\beta} (\alpha\delta + \gamma\beta + 2\beta\delta p) = \frac{\alpha\delta}{\beta} + \gamma + 2\delta p. \quad (13)$$

With the demand curve, AC curve, and MC curve in hand, we can find where they intersect and compute any area of interest between them. In our baseline (linear) specification, the intersection points and areas of interest can be computed using simple geometry. The equilibrium price and quantity are given by equating $AC(p) = D(p)$, resulting in $P_{eq} = \frac{\gamma}{1-\delta}$ and $Q_{eq} = \alpha + \beta \frac{\gamma}{1-\delta}$. The efficient price and quantity are given by equating $MC(p) = D(p)$, resulting in $P_{eff} = \frac{1}{1-2\delta} \left(\frac{\alpha\delta}{\beta} + \gamma \right)$ and $Q_{eff} = \alpha + \frac{1}{1-2\delta} (\alpha\delta + \beta\gamma)$. The efficiency cost associated with competitive pricing (measured by the area of triangle CDE in Figure 1) is then given by

$$\Delta_{CDE} = \frac{1}{2} (Q_{eff} - Q_{eq}) (P_{eq} - MC(P_{eq})) = \frac{-\delta^2}{2(1-2\delta)\beta} \left(\alpha + \frac{\beta\gamma}{1-\delta} \right)^2. \quad (14)$$

In the online appendix we also report results from other, non-linear, specifications, in which we compute these price, quantity, and welfare estimates numerically.

4.4 Baseline results

Our baseline specification estimates the linear demand and cost curves shown in equations (11) and (12) on our baseline sample. This allows us to walk through the main conceptual points of interest involved in applying our proposed approach. In the online appendix we provide a more thorough and detailed discussion of empirical issues specific to our context, including alternative samples and specifications.

Table 2 shows the raw data for our key variables. The (relative) price ranges from \$384 to \$659, with about three-quarters of the sample facing the lowest price. Column (3) shows that the propensity to choose contract H is generally declining with price, and ranges from 0.67 to 0.43. Column (4) shows that the average costs of the (endogenously selected) individuals who select contract H is generally increasing with price (or equivalently, declining in quantity). This pattern of average costs indicates the existence of adverse selection (see Figure 1). Column (5) shows the same for the individuals who (endogenously) select contract L . Recall that incremental cost is defined as the difference in costs to the insurer associated with a given employee’s family’s medical expenditures if those expenditures were insured under contract H rather than contract L . As shown in Figure 3, this difference is a non-linear function of expenditures.

In the spirit of the “positive correlation” test (Chiappori and Salanie 2000), a comparison of columns (5) and (4) reveals consistently higher average costs for those covered by contract H compared to those covered by contract L . This indicates that either moral hazard or adverse selection is present. Detecting whether selection is present, and if so what its welfare consequences are, requires the use of our pricing variation, to which we now turn.

Using OLS, we estimate equation (11) with no additional controls and obtain a downward sloping demand curve, with a slope coefficient β of -0.00070 (with a standard error of 0.00032; p -value = 0.034) and a constant α of 0.94 (standard error of 0.12; p -value < 0.01). This implies that a \$100 increase in price reduces the probability that the employee chooses contract H by a statistically significant 7 percentage points, or about 11%.

Similarly, we use OLS to separately estimate the average cost curve in equation (12). We obtain a slope coefficient δ of 0.155 (standard error of 0.064; p -value = 0.021) and a constant of 391.7 (standard error of 26.8; p -value < 0.01). As noted, the slope of the cost curve represents a test for the existence and nature of selection, and the positive coefficient on price indicates the presence of adverse selection. That is, the average cost of individuals who purchased contract H is (statistically significantly) higher when the price is higher. In other words, when the price selects those who have, on average, higher willingness to pay for contract H , the average costs of this group are also higher. The average cost curve is therefore downward sloping (in quantity, as in Figure 1).

The point estimate from our baseline specification suggests that a dollar increase in the relative price of contract H is associated with an increase in the average cost of the (endogenous) sample selecting contract H at that price of about 16 cents. By itself, this estimate of the cost curve can only provide evidence of the existence of adverse selection. Without knowledge of the demand curve, it does not allow us to form even an approximate guess of the associated efficiency cost of

adverse selection. A central theme of this paper is that we can combine the estimates from the demand curve and the cost curve to move beyond detecting selection to quantifying its efficiency cost and, relatedly, to calculating the welfare benefits from a set of public policy interventions.²⁴

In this spirit, Figure 5 shows how to translate the baseline empirical estimates of the demand and cost curves into the theoretical welfare analysis. That is, Figure 5 presents the empirical analog to Figure 1 by plotting the estimated demand and average cost curves, as well as the marginal cost curve implied by them (see equation (13)). Based on these estimates, it is straightforward to calculate the implied welfare cost of competitive pricing, i.e. area CDE in Figure 5 (and Figure 1). It should be readily apparent from the figure that, holding the cost curve constant, shifting and/or rotating the demand curve could generate very different welfare costs. This underscores the observation that merely estimating the slope of the cost curve is not by itself informative about the likely magnitude of the resultant inefficiency.

We estimate that the welfare cost associated with competitive pricing is \$9.55 per employee per year, with a 95% confidence interval ranging from \$1 to \$40 per employee.²⁵ Adverse selection raises the equilibrium price by almost \$200 above the efficient price (compare the estimated efficient price at point E to the estimated equilibrium price at point C), and correspondingly lowers the share of contract H by 14 percentage points. The social benefit of providing contract H to the marginal employee who buys contract L in equilibrium (i.e. the vertical distance between points C and D in Figure 5) is \$138.

Figure 5 also provides some useful information about the fit of our estimates, and where our pricing variation is relative to the key prices of interest for welfare analysis. The circles superimposed on the figure represent the actual data points (from Table 2), with the size of each circle proportional to the number of individuals who faced that price. The fit of the cost curve appears quite good. The fit of the demand curve is also reasonable, although the scatter of data points led us to assess the sensitivity of the results to a concave demand curve, which is one of the exercises reported in the online appendix. The price range of \$384 to \$659 in our data brackets our estimate of the equilibrium price (point C) of \$463. The lowest (and modal) price in our sample of \$384 is about 45% higher than our estimate of the efficient price (point E) of \$264. Thus, while in principle our approach does not require parametric assumptions, in practice the span of the pricing variation in our particular application requires that we impose some functional form assumptions to estimate the area of triangle CDE. In the online appendix we examine alternative functional forms.

²⁴As noted in Section 2.4, when contract L involves partial coverage, welfare analysis would need to account for the (negative) externalities associated with any moral hazard effects. Our analysis here does not account for such effects since, as we show and discuss in the online appendix, we are unable to reject the null of no moral hazard in our specific application.

²⁵We computed this confidence interval using non-parametric bootstrap. That is, we draw 1,000 bootstrapped samples, and repeated our baseline analysis in each sample. The 95% confidence interval is given by the 2.5th and 97.5th percentiles in the distribution of welfare cost estimates.

4.5 Welfare analyses

We show how our framework can be used to produce a number of other welfare estimates. These may be of interest in their own right and also serve as useful comparisons to our baseline estimate of the welfare cost of inefficient pricing arising from adverse selection (triangle CDE).

Benchmarks for our welfare cost estimates. We can use the demand and cost curves shown in Figure 5 to calculate various benchmarks that provide some context for our estimate of the welfare cost of competitive pricing of \$9.55 per employee. An important consideration in choosing a benchmark is how out of sample we must take the demand and cost estimates in order to form it. Again, Figure 5 is informative on this point.

We calculate two useful denominators to scale our estimate of the welfare cost. One is a measure of how large this cost could have been before we started the analysis. Our thought experiment is to assume that we observe data (on price, quantity, and costs) from only one of the rows of Table 2, so there is no price variation. We assume we observe the weighted average price of \$414. Since individuals have the option to buy contract H at this price but choose not to do so, their welfare loss from not being covered by contract H cannot exceed \$414. Our estimate of the efficiency cost of \$9.55 is therefore 2.3% of this “maximum money at stake,” as Einav, Finkelstein, and Schrimpf (2009) term this construct.

A second useful denominator is to scale the welfare cost from competitive pricing arising from adverse selection by the total surplus at stake from efficient pricing. We therefore calculate the ratio of triangle CDE (the welfare loss from competitive pricing) to triangle ABE (the total welfare from efficient pricing) in Figure 1. To enhance readability, points A and B are not shown in Figure 5, but are easily calculated from the parameter estimates. They are, however, fairly out of sample relative to our data. For example, at point A we estimate price to be about \$1,350, which is more than twice the highest price we observe in the data. In our particular application therefore, this benchmark raises concerns about extrapolating too far out of sample, although we show in the online appendix that the result is relatively robust to alternative functional forms for that extrapolation) Using this benchmark as a denominator, we estimate that the welfare loss from adverse selection is about 3% of the surplus at stake from efficient pricing.

Welfare under other market allocations. Although our welfare analysis has focused on the efficiency cost of competitive equilibrium pricing arising from adverse selection, the fact that we observe prices varying – and this is how we identify the demand and cost curves – underscores the point that to generate our pricing variation we observed a market that is not in equilibrium. Our analysis of “equilibrium” pricing, like our analysis of “efficient” pricing, is based on a counterfactual. By the same token, our analysis of the efficiency cost of such pricing is not an analysis of the realized efficiency cost for our sample but rather what this efficiency cost would be if, contrary to fact, these options were offered in a competitive market setting. Since our demand and cost curves are sufficient statistics for welfare analysis of the pricing of existing contracts, we can use them to compute the welfare cost of any other inefficient pricing. For example, we estimate that the weighted average of

the welfare cost of adverse selection given the observed pricing in our sample (see Table 2, columns (1) and (2)) is \$6.26 per employee per year.

Moreover, as we noted in Section 2, we could also use the estimated demand and cost curves to estimate welfare under alternative assumptions about the market equilibrium, including monopoly or imperfect competition. For example, a monopolist facing our estimated demand and cost curves would set a (relative) price of \$907 for contract H . The resultant efficiency cost would be just below \$100 per employee, which is an order of magnitude higher than the estimated efficiency cost from competitive pricing.

Another interesting alternative is to compute what the welfare cost of competitive pricing would be if, contrary to what happens in the employment context, competitive prices were set based on some observable characteristics of the employees. To do so, we simply estimate the demand and cost curves separately for each “cell” of individuals who, based on their characteristics, would be offered the same price. As an example, we consider what would happen to our welfare estimate if prices were set differently based on whether the family coverage applied to 3 individuals, 4 individuals, or 5 or more individuals. About half of our baseline sample has 4 covered members, and the remaining sample is evenly split between the other two categories. We maintain the assumption that the equilibrium would involve average cost pricing, although now the equilibrium is determined separately in each of the three market segments. We detect adverse selection in each segment separately, and estimate that the (weighted average) welfare cost of this selection would be \$12.92 if prices were set differently for each market segment, compared to our estimated welfare cost of \$9.55 when family size is not priced.

Welfare consequences of government intervention. Adverse selection provides the textbook economic rationale for government intervention in insurance markets. We therefore show how we can use our framework to estimate the welfare cost of standard public policy interventions in insurance markets. We then compare this to our estimate of the welfare cost of competitive pricing. As mentioned, our approach allows us to analyze the welfare consequences of counterfactual public policies that change the price of existing contracts, such as price subsidies, coverage mandates, and regulation of the characteristics of individuals that can be used in pricing. This last potential policy was already discussed in the previous section where we analyzed the welfare consequences of firms pricing on a characteristic (in our example, family size) that is not currently priced.

Our preferred policy analysis in our particular application is to compare the social welfare gain from efficient pricing (triangle CDE) to the social welfare cost of the price subsidy required to achieve this efficient price. An attraction of this calculation is that it does not require further out of sample extrapolation beyond what is needed to compute the area of triangle CDE itself. The social cost of such a subsidy is given by $\lambda(P_{eqm} - P_{eff})Q_{eff}$ where λ is the marginal cost of public funds. Given our estimates of the efficient and equilibrium outcomes (Figure 5), and using 0.3 as the (standard estimate of) marginal cost of public funds (e.g., Poterba 1996), we calculate the social cost of the price subsidy needed to achieve the efficient allocation to be \$45. That is, we estimate that the social cost of a price subsidy that achieves the efficient allocation is about five

time larger than the social welfare (of \$9.55) it gains.

Of course, given a non-zero social cost of public funds, the welfare maximizing subsidy would not attempt to achieve the efficient allocation. It is therefore also interesting to investigate whether there is any scope for welfare improving government intervention in the form of a price subsidy to contract H . We do this by investigating whether, at the competitive allocation (point C), a marginal dollar of subsidy is welfare enhancing. We calculate that in our application it is not, so that the welfare maximizing (additional) price subsidy by the government is therefore zero.²⁶

We also compared welfare in the competitive equilibrium with adverse selection to welfare when everyone is mandated to be covered by contract H . Mandatory insurance is the canonical solution to the problem of adverse selection in insurance markets (Akerlof 1970), making the analysis of the mandate of considerable interest.²⁷ However, in our application, the welfare cost of mandating coverage by contract H (area EGH in Figure 1) requires calculating points which are reasonably far out of sample. This suggests that in our particular application more caution is warranted with this analysis (although again we show in the online appendix that the estimate is reasonably robust). With this important caveat in mind, we estimate that the welfare cost from mandatory coverage by contract H is about three times higher than the welfare cost associated with competitive pricing.

5 Conclusions

This paper proposed a simple approach to quantifying and estimating the welfare cost caused by inefficient pricing in insurance markets with selection. We show how standard consumer and producer theory can be applied to welfare analysis of such markets, and we provide a graphical representation of the efficiency cost of competitive pricing. This graphical analysis not only provides helpful intuition but also suggests a straightforward empirical approach to welfare analysis. The key to estimation is the existence of identifying variation in the price of insurance. Applied welfare analysis usually requires pricing variation that allows the researcher to trace out a demand curve. The defining feature of selection markets is that costs vary endogenously as market participants respond to price. Welfare analysis in such markets therefore requires that we also trace out the (endogenous) cost curve. We show that this is straightforward to do using direct data on cost and the same price variation used to identify demand. In doing so, the slope of the estimated cost curve also provides a direct test of the existence and nature of selection.

We illustrated our framework by applying it in the context of employer-provided health insurance at a particular firm. We find evidence of adverse selection in the market, but we estimate that the welfare cost of the resultant inefficient pricing is quantitatively small. It is important to emphasize

²⁶The marginal benefit from the first dollar of subsidy is \$137.4 (the distance between point C and point E) times the marginal number of newly covered individuals (0.0007 given our estimates of the demand curve). By contrast, the marginal cost of the dollar subsidy is the cost of public funds (0.3) times all of the inframarginal individuals at point C (i.e. 0.617).

²⁷Footnote 5 discussed some of the possible factors that may make it inefficient to allocate the H contract to the entire market.

that our empirical estimates are specific to our particular setting and there is no reason to think that our welfare estimates are representative of other populations, other institutional environments, or other insurance markets. However, at a broader level, our findings illustrate that it is empirically possible to find markets in which adverse selection exists and impairs market efficiency, but where the efficiency cost of the pricing it produces may not be large, or obviously remediable using standard public policy tools. Whether the same is true in other markets and in which is an important area for future work.

We hope that such future work will apply our framework and strategy to other insurance settings (or, more generally, to other settings with hidden information such as credit markets or regulated monopolists). The approach is relatively straightforward to implement and fairly general. As a result, comparisons of welfare estimates obtained by this approach across different settings may be informative. In any given application, we see the transparency of our approach as one of its key attractions. The direct mapping from the theoretical framework (Figure 1) to its empirical analog (Figure 5) facilitates an informed appraisal of the estimates, including such issues as in-sample fit, the extent of out-of-sample extrapolation needed for a particular welfare estimate, and the extent and validity of the pricing variation.

As we emphasize throughout, our approach is unable to shed light on the welfare consequences of any distortion in the contract space induced by selection, or of public policies that introduce contracts not observed in the data. Analysis of such questions would require a model of the primitives underlying the revealed demand and cost curves. We view such models as a useful and important complement to the empirical approach we have proposed here.

6 References

Abbring, Jaap, Pierre-Andre Chiappori, and Jean Pinquet, "Moral hazard and Dynamic Insurance Data," *Journal of the European Economic Association*, 1 (2003), 767-820.

Adams, William, Liran Einav, and Jonathan Levin, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99 (2009), 49-84.

Akerlof, George, "The Market for 'Lemons': Quality Uncertainty and The Market Mechanism," *Quarterly Journal of Economics*, 84 (1970), 488-500.

Barseghyan, Levon, Jeffrey Prince, and Joshua Teitelbaum, "Are Risk Preferences Stable across Contexts? Evidence from Insurance Data," mimeo, Cornell University, 2008.

Blackmon, Glenn B., and Richard Zeckhauser, "Mispriced Equity: Regulated Rates for Automobile Insurance in Massachusetts," *American Economic Review Papers and Proceedings*, 81 (1991), 65-69.

Buchmueller, Thomas C., and John E. DiNardo, "Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania, and Connecticut," *American Economic Review*, 91 (2002), 280-294.

Bundorf, Kate M., Jonathan Levin, and Neale Mahoney, "Pricing, Matching and Efficiency in Health Plan Choice," mimeo, Stanford University, 2008.

- Bundorf, Kate M., and Kosali Simon, "The Impact of Rate Regulation on Access to Supplemental Health Insurance," *American Economic Review Papers and Proceedings*, 96(2006), 67-71.
- Busch, Susan H., Colleen L. Barry, Sally J. Vegso, Jody L. Sindelar, and Mark R. Cullen, "Effects of a Cost-Sharing Exemption on Use of Preventive Services At One Large Employer," *Health Affairs*, 25 (2006), 1529-1536.
- Carlin, Caroline, and Robert J. Town, "Adverse Selection, Welfare and Optimal Pricing of Employer-Sponsored Health Plans," mimeo, University of Minnesota, 2007.
- Chetty, Raj, "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116 (2008), 173-234.
- , "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced Form Estimates," *Annual Review of Economics*, 1 (2009), forthcoming.
- Chetty, Raj, and Emmanuel Saez, "Optimal Taxation and Social Insurance with Endogenous Private Insurance," NBER Working Paper 14403, 2008.
- Chiappori, Pierre-Andre, and Bernard Salanie, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108 (2000), 56-78.
- Cohen, Alma, and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97 (2007), 745-788.
- Cutler, David M., Amy Finkelstein, and Kathleen McGarry, "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," *American Economic Review Papers and Proceedings*, 98 (2000), 157-162.
- Cutler, David M., and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *Quarterly Journal of Economics*, 113 (1998), 433-466.
- Cutler, David M., and Richard J. Zeckhauser, "The Anatomy of Health Insurance," in *Handbook of Health Economics*, A. J. Culyer and J. P. Newhouse, eds., Volume 1, 563-643 (Elsevier, 2000).
- De Meza, David, and David C. Webb, "Advantageous Selection in Insurance Markets," *Rand Journal of Economics*, 32 (2001), 249-262.
- Eichner, Matthew J., Mark B. McLellan, and David A. Wise, "Insurance or Self-Insurance? Variation, Persistence, and Individual Health Accounts," in *Inquiries in the Economics of Aging*, David Wise, ed. (Chicago, IL: University of Chicago Press, 1998).
- Einav, Liran, Amy Finkelstein, and Jonathan Levin, "Beyond Testing: Empirical Models of Insurance Markets," *Annual Review of Economics*, 2 (2010), forthcoming.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf, "Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market," mimeo, MIT and Stanford University, 2009.
- Finkelstein, Amy, and Kathleen McGarry, "Multiple Dimensions of Private Information: Evidence from The Long-Term Care Insurance Market," *American Economic Review*, 96 (2006), 938-958.
- Finkelstein, Amy, and James Poterba, "Selection Effects in the Market for Individual Annuities: New Evidence from the United Kingdom," *The Economic Journal*, 112 (2002), 28-50.

———, “Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market,” *Journal of Political Economy*, 112 (2004), 193-208.

Fox, Peter, Thomas Rice, and Lisa Alexih, “Medigap Regulation: Lessons for Health Care Reform,” *Journal of Health Politics, Policy and Law*, 20 (1995), 31-48.

Greenwald, Bruce C., and Joseph E. Stiglitz, “Externalities in Economies with Imperfect Information and Incomplete Markets,” *Quarterly Journal of Economics*, 101 (1986), 229-264.

Gruber, Jonathan, “The Incidence of Mandated Maternity Benefits,” *American Economic Review*, 84 (1994), 622-641.

———, “Taxes and Health Insurance,” in *Tax Policy and the Economy*, James Poterba, ed., Volume 16, 37-16 (Cambridge, MA: MIT Press, 2002).

Israel, Mark, “Do We Drive More Safely When Accidents Are More Expensive? Identifying Moral Hazard from Experience Rating Schemes,” unpublished mimeo, available at http://www.wcas.northwestern.edu/csio/Conferences/CSIO-IDEI-2004/moral_exp_paper1.pdf, 2004.

Karlan, Dean, and Jonathan Zinman, “Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment,” *Econometrica*, 2009, Forthcoming.

Keating, Bryan, “Switching in an Evolving Industry,” Ph.D. Dissertation, Stanford University, Chapter 1, 2007.

Lustig, Joshua D., “The Welfare Effects of Adverse Selection in Privatized Medicare,” mimeo, Boston University, 2008.


Poterba, James, “Government Intervention in the Markets for Education and Health Care: How and Why?” in *Individual and Social Responsibility*, Victor Fuchs, ed. (Chicago, IL: University of Chicago Press, 1996).

Rothschild, Michael, and Joseph E. Stiglitz, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90 (1976), 630-649.

Stiglitz, Joseph E., “The Causes and Consequences of the Dependence of Quality on Price,” *Journal of Economics Literature*, 25 (1987), 1-48.


Sydnor, Justin, “Sweating the Small Stuff: The Demand for Low Deductibles in Homeowners Insurance,” mimeo, Case Western University, 2006.

Figure 1: Efficiency cost of adverse selection




This figure represents the theoretical efficiency cost of adverse selection. It depicts a situation of adverse selection because the marginal cost curve is downward sloping (i.e. increasing in price, decreasing in quantity), indicating that the people who have the highest willingness to pay also have the highest expected cost to the insurer. Competitive equilibrium is given by point C (where the demand curves intersects the average cost curve), while the efficient allocation is given by point E (where the demand curve intersects the marginal cost curve). The (shaded) triangle CDE represents the welfare cost from under-insurance due to adverse selection.

Figure 2: Efficiency cost of advantageous selection




This figure represents the theoretical efficiency cost of advantageous selection. It depicts a situation of advantageous selection because the marginal cost curve is upward sloping, indicating that the people who have the highest willingness to pay have the lowest expected cost to the insurer. Competitive equilibrium is given by point C (where the demand curve intersects the average cost curve), while the efficient allocation is given by point E (where the demand curve intersects the marginal cost curve). The (shaded) triangle CDE represents the welfare cost from over-insurance due to advantageous selection.

Figure 3: Description of contract H and contract L




Figures 3(a) and 3(b) present the main features of contract H (dashed) and contract L (solid) family coverages offered by the company, which is based on a deductible and an out-of-pocket maximum. Figures 3(c) and 3(d) present the corresponding cost difference to the insurer by providing the contract H instead of contract L , for a given level of medical expenditure. In other words, Figures 3(c) and 3(d) illustrate the in-network and out-of-network components of the constructed variable $c_i(m)$. Figures 3(a) and 3(c) describe the rules for in-network medical spending (deductibles of \$0 and \$500, and out-of-pocket maximums of \$5,000 and \$5,500 for contracts H and L , respectively), and Figures 3(b) and 3(d) describe the rules for out-of-network medical spending (deductibles of \$500 and \$1000, and out-of-pocket maximums of \$10,000 and \$11,000 for contracts H and L , respectively). Coinsurance rates for both contracts are 10% (in network) and 30% (out of network). There is no interaction between the in-network and out-of-network coverages (i.e. each deductible and out-of-pocket maximum must be satisfied separately). The online appendix provides more details on the coverage rules and our construction of $c_i(m)$.

Figure 4: The distribution of the insurer's incremental costs (c_i)



This figure presents the distribution of the incremental insurer cost (c_i) for all 3,779 employees in our baseline sample. Note that the distribution has several mass points which are driven by the kinked formula of the coverages (Figure 3). The largest mass point is at \$450, with about two thirds of the sample. This point represents individuals who spent more than \$500 and less than \$50,000 in network, and less than \$500 out of network.

Figure 5: Efficiency cost of adverse selection - empirical analog



This figure is the empirical analog of the theoretical Figure 1. The demand curve and AC curve are graphed using the point estimates of our baseline specification reported in the text. The MC curve is implied by the other two curves, as in equation (13). The circles represent the actual data points (Table 2) for demand (empty circles) and cost (filled circles). The size of each circle is proportional to the number of individuals associated with it. For readability we omit the one data point from Table 2 with only 7 observations (although it is included in the estimation). We label points C, D, and E, that correspond to the theoretical analog in Figure 1, and report some important implied point estimates (of the equilibrium and efficient points, as well as the welfare cost of adverse selection).

Table 1: Assessing the exogeneity of the price variation

	Faced lowest relative price (2,939 employees) (1)	Faced higher relative prices (840 employees) (2)	Difference (3)	Coefficient (4)	p-value (5)
Age (Mean)	42.74	42.40	0.33	-0.245	0.31
Tenure (Mean)	13.02	11.63	1.39	-0.565	0.08
Fraction Male	0.862	0.852	0.009	1.268	0.79
Fraction White	0.874	0.825	0.049	-6.998	0.40
Log(Annual Salary) (Mean)	11.16	11.05	0.11	-8.612	0.17
Spouse Age (Mean)	41.37	41.05	0.32	-0.200	0.41
Number of covered family members (Mean)	4.14	4.07	0.07	-1.400	0.36
Age of youngest covered child (Mean)	9.81	9.41	0.40	-0.3	0.26
Log(2003 Medical Spending + 1) ^a					
All	8.13	7.79	0.32	-2.100	0.15
In most common 2003 plan	8.21	8.08	0.13	-1.700	0.08

The table reports average differences in covariates (shown in the left column) across employees who face different relative prices for the higher coverage option in the baseline sample. The employee characteristics in the left column represent contemporaneous 2004 characteristics (except where noted). Note that everyone with family coverage has a covered spouse and at least one covered child. Columns (1) and (2) present, respectively, average characteristics for the approximately three-quarters of employees who faced the lowest relative price (\$384; see Table 2) and the remaining one quarter who face one of the five higher relative prices (\$466 to \$659; see Table 2). Column (3) shows the difference between columns (1) and (2). Columns (4) and (5) report, respectively, the coefficient and p-value from a regression of the (continuous) relative price variable (in \$US) on the characteristic given in the left column; we adjust the standard errors for an arbitrary variance covariance matrix within each state.

^a In the bottom two rows we look at 2003 medical spending for all employees in the sample who were in the data in 2003 (2,602 and 658 employees in columns (1) and (2), respectively), and for all employees who were in the data in 2003 in the most common 2003 health insurance plan (2,284 and 523 employees in columns (1) and (2), respectively). The latter attempts to avoid potential differences in spending arising from moral hazard effects of different 2003 coverages.

Table 2: The effect of price on demand and costs

(Relative) Price	Number of employees	Fraction chose contract H	Average incremental cost for those covered under:	
			Contract H	Contract L
(1)	(2)	(3)	(4)	(5)
\$384	2,939	0.67	\$451.40	\$425.48
\$466	67	0.66	\$499.32	\$423.30
\$489	7	0.43	\$661.27	\$517.00
\$495	526	0.64	\$458.60	\$421.42
\$570	199	0.46	\$492.59	\$438.83
\$659	41	0.49	\$489.05	\$448.50

The table presents the raw data underlying our baseline estimates. All individuals face one of six different (relative) prices, each represented by a row in the table. Column (2) reports the number of employees facing each price, and column (3) reports the fraction of them who chose contract H . Columns (4) and (5) report (for individuals covered by contracts H and L , respectively) the average incremental costs to the insurer of covering these individuals with contract H rather than with contract L , taking the family's medical expenditures as given. The graphical analog to this table is presented by the circles shown in Figure 5.

Online Appendix

“Estimating welfare in insurance markets using variation in prices” by
Einav, Finkelstein, and Cullen

A.1. Data construction

Construction of the baseline sample. We make a number of sample restrictions. First, we make a number of restrictions for purposes of data purity, which brings the original sample of about 45,000 active employees down to about 37,000 active employees. The biggest reduction in sample size comes from excluding employees who are not at the company for the entire year (for whom we do not observe complete annual medical expenditures, which are necessary for estimating the cost curve). In addition, we exclude employees who are outside the traditional benefit structure of the company (for example, because they were working for a recently acquired company with a different (grandfathered) benefit structure). For such employees we do not have detailed information on their insurance options and choices. We also exclude a small number of employees because of missing data or data discrepancies.

Second, because the new set of health insurance options we study did not apply to many hourly employees and because (as we discuss in Section 4.2) the pricing variation is cleaner for the salaried employees, we further limit the analysis in this paper to salaried employees, who are approximately one third of the U.S.-based Alcoa workforce.

Third, to illustrate most easily how the theoretical framework maps to the empirical strategy, we limit the baseline analysis to the two modal health insurance choices: a higher and a lower level of PPO coverage, to which we refer as contract H and contract L throughout Section 4 of the paper. Approximately two-thirds of salaried employees chose one of these two PPO options. In Section A.2 of this online appendix we describe the other options in more detail and show that the pricing of the two PPO options we focus on does not affect the probability of the employee choosing one of the other options we do not analyze. This helps to alleviate concerns about potential biases from our sample selection on contract choice.

Finally, for simplicity, our baseline specification further limits our sample of salaried employees who choose either contract H or contract L to the slightly over one half of such employees who choose the most common coverage tier, which is family coverage. All employees have a choice of four different tiers for health insurance coverage: employee only, employee plus spouse, employee plus children, and family coverage. In Section A.2 of this online appendix we show that our results are similar when we include employees in all coverage tiers. We assume throughout that the choice of coverage tier is unrelated to the pricing variation. A priori, this seems a reasonable assumption given that coverage tier options are limited by the demographic composition of the family, and that the price multiplier across coverage tiers is the same for all employees.¹ Consistent with our

¹Specifically, for any health insurance coverage option, for all employees the family price is always triple the

assumption, we find that the (relative) price of contract H in the family coverage tier does not predict (either economically or statistically) which coverage tier the employee chooses (not shown).

Table A1 provides some descriptive statistics on the employees. Column (1) presents descriptive statistics for the sample of 37,000 active employees for whom we have complete data. Column (2) limits the sample to the approximately one third of the sample who are salaried employees. Column (3) makes the further (minimal) restriction to the salaried employees who face the new benefit design. Column (4) further limits the sample to employees who choose either contract H or contract L , and column (5) further limits the sample to those in family coverage. Column (5) represents our baseline sample that we use for most of the empirical analysis. Section A.2 of this online appendix presents analyses that use all coverage tiers (column (4)) and all coverage options (column (3)).

For comparison, columns (6) through (8) of Table A1 present statistics from the 2005 March Current Population Survey (CPS) on characteristics of various types of full-time employees in the U.S. The principal (and unsurprising) finding is that Alcoa employees do not appear to be representative of any cut of full time employees in the U.S. We also compared the medical expenditures in our baseline sample to medical expenditure data from the 2004 Medical Expenditure Panel Survey (MEPS). Salaried employees in Alcoa tend to have about 50 percent lower medical expenditures than comparable individuals in the MEPS.² This may be because Alcoa salaried employees are healthier than the general population or that they tend to live in regions with lower healthcare costs.³ Such comparisons underscore our statements that our empirical results should be viewed as an illustrative example of how our proposed approach can be applied, rather than as generalizable findings about the employer-provided health insurance market in the U.S.

Construction of incremental costs. One of the key variables in our analysis is the insurer’s incremental costs c_i . This is defined as $c(m_i; H) - c(m_i; L)$, where $c(m_i; j)$ is the cost to the insurer from covering medical expenditures m_i under contract j . We note that medical utilization (m_i) is held fixed in the construction of c_i , so there is no estimation involved in the process.

The construction of c_i requires detailed knowledge of each plan’s benefits as well as individuals’ realized medical claims. We obtained the former from reading each plan coverage details and verifying them with the actual reimbursements we observe in the data. The latter is part of our data, which include detailed information about every single claim made by Alcoa employees during 2004. For each claim we know the claim date, the claim amount, how much of it was reimbursed by Alcoa, and how much was paid out of the insured’s pocket. For the latter we also know whether it was applied to the annual deductible or was part of a coinsurance. We also know whether each

“employee only” price, 1.58 times the “employee plus children” price, and 1.43 times the “employee plus spouse” price.

²Specifically we focus on MEPS observations on individuals with full-year coverage by employer-provided health insurance, and we try to reweight observations to adjust for age and gender.

³In addition to the non-representativeness of health expenditures in our Alcoa population, we further note that our cost variable is a complicated non-linear transformation of total cost, which is perhaps even more context-specific (as the transformation depends on the particular features of the plans we study).

claim was associated with in-network or out-of-network care, and additional medical details which are less relevant for the construction of c_i .

To construct c_i we simply “run” the observed set of claims for each covered employee (and his dependents) through the reimbursement rules twice. Once by applying the rules of contract H and once by applying the rules of contract L . A key feature of our setting which facilitates this construction is that the two contracts we focus on vary only in their employee cost-sharing rules. Alcoa is the direct insurer of both plans, and the plans are identical in all other features, such as the network definition and the benefits covered. As a result, we do not have to worry about differences between contracts H and L in plan features that might differ in unobservable ways across employees (for example, differences in providers or provider prices, the relative network quality, and so forth). In particular, this implies that a set of claims submitted under one contract would be eligible (and identical) claims under the other contract. Once we have “run” the claims for each employee for each contract, we have obtained $c(m_i; H)$ and $c(m_i; L)$, and the difference is our constructed variable c_i .

Applying the plans’ rules is fairly simple, although certain issues require some care. One such issue is whether a claim was made in network or out of network, since different deductible and cost sharing rules would apply (see Figure 3). A second issue is related to preventive care. Alcoa provides full coverage (with zero out-of-pocket payments) for various preventive treatments, including periodical exams, well baby, etc. It is therefore important to know whether claims are associated with preventive-related services, since cost sharing rules do not apply to such claims. A third issue, which is typical of most health insurance plans, is the interaction between an individual deductible and a family deductible (as well as analogous issues regarding individual and family out-of-pocket maximums). In our data, the family deductible is always twice the individual deductible. For a family with more than two covered individuals, it is therefore important to account for the interaction among family members, as the cost sharing rules would vary depending on how the spending is distributed among the family members. That is, a given individual in a family can exhaust his deductible either by spending his individual deductible or by having the cumulative spending of other members of the family reach the family deductible. In the construction of c_i , we therefore need to account for the composition of spending within the family.

Fortunately, the data are quite detailed and the plan rules are fairly simple (despite the above issues), allowing us to calculate $c(m_i; j)$ with a great deal of accuracy. Indeed, our calculated reimbursements (based on our application of the plan rules) and the actual reimbursements observed in the data are almost the same. For example, for individuals with contract H the correlation between their actual (observed) share of out-of-pocket spending (out of total expenditure) and our constructed share is over 0.97. The same is true for contract L , or when we correlate levels of expenditures instead of shares. Recall that we still need to apply our construction, because for each individual we only observe the actual reimbursement for the contract he chose, while the second element of c_i is always a counterfactual. For consistency, we never use the actual reimbursement and always compute c_i by constructing both elements, $c(m_i; H)$ and $c(m_i; L)$.

A.2. Robustness of the baseline estimates

In this section we explore the sensitivity of our welfare estimates to a number of alternative specifications. Our overall finding is that the magnitude of the various welfare estimates discussed in the paper – even those that involve extrapolation considerably out of sample – are qualitatively similar across a range of alternative specifications. In particular, across various alternative specifications, the welfare gain from a price subsidy that achieves the efficient price is always substantially below the social cost of the required price subsidy. Similarly, the welfare loss from competitive pricing when choice over contracts is allowed is always lower than the welfare loss from mandatory coverage by contract H , and the welfare cost of competitive pricing is always less than 10 percent of the total surplus that could be generated from efficient pricing. In the end of this section we also address possible concerns regarding sample selection.

Functional form and theoretical restrictions on the demand curve. Table A2 summarizes some of the sensitivity analyses. Panel A summarizes the implied welfare implications of each specification. For completeness, Panel B shows the corresponding parameter estimates from each specification (which are used to derive the welfare estimates shown in Panel A). In the interest of brevity we focus our discussion primarily on the robustness of the resultant welfare estimates (columns (6) through (8) of Panel A), which are our main interest. The first row of Table 5 presents the results from our baseline specification discussed in the main text. Subsequent rows reports results from a single, specified departure from this baseline.

Rows 2-5 in Table A2 explore the sensitivity of our results to our functional form assumptions. Row 2 shows the results from our baseline specification are quite similar if we estimate a probit for the demand equation rather than a linear demand. Unrestricted quadratic demand (not reported) behaves very badly out of sample and is therefore not shown (but in row 5 we report and discuss a restricted specification that includes a quadratic demand curve). As can be seen in Figure 5, the linear specification fits the cost data well.⁴

We also experimented with imposing restrictions on the demand curve that are implied by basic price theory. Willingness to pay is (theoretically) bounded from above at \$1,500 (the maximum possible out-of-pocket savings from contract H ; see Figure 3) and (theoretically) bounded from below by 0 (any rational individual should always prefer more coverage to less if the former is offered for free). Our baseline demand estimate (Table A2, row 1) satisfies the first constraint (the share of contract H becomes 0 at a price of \$1,350), but not the second. At a price of 0, the share of contract H is only 0.94.⁵ The results in row 3 show that constraining the share of contract H to be 1 when price is 0 does not noticeably affect our welfare estimates. Row 4 shows the results are

⁴We explored alternative functional forms for the cost curve, such as a quadratic, log-log, and log-linear functions. Not surprisingly, the results (not shown) were very similar in sample. However, curvature (concavity in particular) in the estimated AC curve sometimes led to out-of-sample predictions that were difficult to interpret (such as non-monotone MC curve). Given all these hard-to-interpret predictions were driven by out-of-sample predictions from an ad hoc functional-form extrapolation, we prefer to simply reject such extrapolations and focus our discussion on those extrapolation that seem to “better behave” (out of sample).

⁵One reason why we may estimate demand below 1 for a price of 0 is that our functional form assumption of

also similar if we impose the constraint that willingness to pay is bounded at \$800, which may be a more reasonable upper bound in practice than the theoretically possible \$1,500.⁶ Row 5 estimates a quadratic demand curve, imposing both the (1,\$0) and the (0,\$800) constraints on (Q,P), and again resulting in welfare estimates that are quite stable.

Tax treatment of employee premiums. We also considered the sensitivity of our results to the tax treatment of employee contributions to health insurance and to out-of-pocket medical expenditures. Employee premium contributions are made pre-tax. Employees can pay their out-of-pocket medical spending pre-tax as well, by contributing to a Flexible Spending Account (FSA). If all out-of-pocket expenses were paid pre-tax, the tax treatment of employee premiums and employee medical spending would be symmetric, and ignoring the tax subsidy to employee premiums (as we do in our baseline specification) would be appropriate. However, in practice, less than a quarter of Alcoa employees contribute to an FSA. It is of course unclear whether employees who do not take advantage of the tax subsidy to out-of-pocket medical spending offered by FSAs are cognizant of the tax subsidy to employee premiums. However, to investigate the sensitivity of our findings to the tax subsidy, we consider the effect on our estimates of assuming that all employees (including those who contribute to FSAs) make their health insurance choices based on the pre-tax price.⁷ We calculate the average tax subsidy (i.e. one minus the average marginal tax rate) for our sample to be 65 percent.⁸ In row 6 we therefore re-estimate the baseline specification with the price variable in both the demand and cost equations multiplied by 0.65. Once again the core welfare estimates are

linear demand is not appropriate for extrapolating this far out of sample. Another possible explanation may be that contract L was the default option in 2004. We suspect that default may be less important in our setting than in others because 2004 was the first year in which the new benefits were offered. These new benefits came with much effort by Alcoa to advertise and explain the new options to its employees, making it likely that most individuals were “active” choosers. Moreover, it is possible to have a model of defaults in which our welfare analysis is unaffected. We discuss this in a little more detail below.

⁶\$1,500 out-of-pocket savings from contract H is only possible if the covered family members spend enough in-network *and* out-of-network to hit the (separate) out-of-pocket maximums. In practice, this never occurs. Indeed, none of the employees in our sample hits the out-of-pocket maximum out-of-network and only about 1 percent hits the in-network out-of-pocket maximum. A potentially more reasonable constraint therefore is that willingness to pay for contract H should not exceed \$800, which is the reduction in out-of-pocket expenditures associated with contract H if the family spends more than the deductible in-network and more than the deductible out-of-network but less than the amount that would cause them to hit the out-of-pocket maximum (see Figure 3).

⁷We do not observe in the data which individuals participate in the FSA.

⁸The tax subsidy is given by $(1 - \tau_f - \tau_s - \tau_{ss} - \tau_{mcr})$ where τ_f is the federal marginal tax rate, τ_s is the state marginal tax rate, τ_{ss} is the marginal Social Security (FICA) payroll tax on the employee, and τ_{mcr} is the marginal Medicare payroll tax on the employee. We estimate these marginal tax rates using the NBER’s TAXSIM model, which takes as inputs the major determinants of marginal tax rates and computes the various marginal rates just mentioned. Many of the required data elements (or reasonable proxies for them) are available in our company’s data, including annual wage and salary income, state, marital status, number of dependents and ages of family members. We assume all employees with family coverage file jointly and do not itemize. We impute wage and salary income of spouse, property income, and dividend income based on the ratio of each of these variables to own income for the sample of full time, white collar manufacturing employees in the March CPS; we pool the 2004-2007 March CPS to increase sample size (Table A1, column (8) presents descriptive statistics for this sample in the March 2005 CPS). All other inputs required by TAXSIM are assumed to be zero. For more information on TAXSIM, see www.nber.org/taxsim.

not noticeably affected, although naturally our estimates of the equilibrium and efficient allocations (see columns (1) through (4)) shift considerably.

Additional covariates and alternative samples. Our baseline estimates of the demand and cost curves include no covariates in the analysis besides the (relative) price. Only variables that are priced should be controlled for in our analysis of selection and its welfare costs. The fact that, for example, individuals of, say, different incomes or different ages may have different expected medical costs, and that this may affect which plan they choose, is part of the endogenous selection we wish to study, rather than control for, since these characteristics are not priced. However, to allow for the possibility that the price menu may be selected differently across states in a systematic fashion (e.g., reflecting differences in healthcare costs across states), in row 7 we include state fixed effects in the demand and cost estimates. Although our estimates become somewhat less precise (see Panel B of Table A2), the welfare implications remain quantitatively similar (Panel A). In row 8 we add all of the contemporary employee characteristics (see Table 1) as covariates to the demand and cost curves (in addition to the state fixed effects).⁹ Once again the results are similar. The fact that the slope of the estimated demand curve remains similar is unsurprising given the evidence in Table 1 that pricing is orthogonal to these employee characteristics. The fact that the slope of the estimated cost curve remains similar suggests that the adverse selection we detect is not driven by the fact that in our setting the observable characteristics of employees are not priced.¹⁰

Finally, in row 9 we estimate our baseline specification using all four coverage tiers rather than just employees with family coverage. Since prices vary by coverage tier, we include (de-meaned) indicator variables for the coverage tier in both the demand and cost estimates.¹¹ The parameter estimates and welfare implications are quite similar to our baseline results.

We also tried restricting our baseline sample, specifically by excluding the 199 individuals who face the \$570 (relative) price, which seem likely to affect the demand estimates. Indeed, we found that eliminating this points substantially reduces the demand elasticity (by about 45%) and it is no longer statistically significant. However, when we do so the average cost curve remains similar. Thus, the steeper demand curve produces a steeper marginal cost curve, exacerbating the welfare costs of inefficient pricing due to selection. As a result, despite the steeper demand curve (which all else equal should reduce welfare costs), our welfare estimate remains roughly the same (\$9.77 compared to \$9.55 in the baseline specification). This type of robustness exercise illustrates that it is the combination of the demand and cost curves that together contribute to the magnitude of

⁹In both rows 7 and 8 the covariates are demeaned so that the constant term is comparable across specifications.

¹⁰In this sense, the robustness test on the cost curve is one sided. Had we found that the slope of the cost curve changed once we controlled for unpriced observables, this would not necessarily be a cause for concern. It could simply reflect the fact that much of the selection in our setting is driven by these unpriced observables.

¹¹The price variable is defined for the chosen coverage tier. As noted earlier, for all employees the prices of contracts in the other coverage tiers are always the same fixed multiplier of the prices in the family coverage tier. To account for the fact that for “employee only” coverage the deductible and out-of-pocket maximum is half of what it is for the other three coverage tiers, we multiply price (p_i) and cost (c_i) by two for the 16% of employees with “employee only” coverage.

the welfare loss.

Possible sample selection. An important potential concern with all of the foregoing analyses is that we limit the sample to only those who choose contract H or contract L , and exclude the approximately one-third of salaried employees who chose one of the five other available options. These five other options are an HMO (chosen by about 7% of salaried employees), opting out of any employer-provided coverage (about 8%), two even lower coverage PPO options (3% in the two of them combined), and a Health Reimbursement Account (HRA) PPO option, which combines a high deductible health insurance policy with tax preferred employer contributions that can be used to pay out-of-pocket expenses (approximately 17%).¹²

In practice, however, our analysis suggests that our sample selection is unlikely to have important effects on our demand estimates (and, of course, it is irrelevant for the estimate of the cost curve which by design is run on the endogenously selected sample of individuals choosing contract H). In particular, we found that the price of contract H relative to contract L (our key right-hand-side variable) does not predict whether or not the employee “opts in” to one of the two contracts we study (contract H and contract L), as opposed to “opting out” into one of the remaining options. We suspect that this in part reflects the fact that many of the other options (in particular the three with non-trivial market share, the HMO, opting out of insurance, and the HRA) are quite horizontally differentiated.

Table A3 presents some of these findings. The dependent variable in the reported linear regressions is an indicator variable that takes the value of 1 if the employee chose one of the “outside goods” and 0 if he chose either contract H or contract L . The right-hand-side variable p is (as before) the relative price of contract H compared to contract L . Column (1) reports the results for employees with family coverage. We find that a \$100 increase in the (relative) price of contract H is associated with an economically and statistically insignificant decline (of 0.09 percentage points) in the probability of choosing one of the outside goods. Column (2) shows similar results when all coverage tiers are pooled. A complication with both of these analyses is that because coverage tier is not available for the 8% of the sample who opt out of coverage, these employees are excluded from the analysis. In column (3) therefore we include in the sample the employees who opt out of coverage. However, since coverage tier is not known for these employees we cannot control for coverage tier and, moreover, we can no longer define the price variable based on the coverage tier. We instead assign all employees the family prices regardless of what coverage tier they actually chose (if known).¹³ Once again there is no evidence that the relative price of contract H has an

¹²The in-network deductibles for the two lower coverage PPO options are \$1,000 and \$1,500. The high deductible HRA PPO has a \$3,000 in-network deductible, but the employee receives \$1,250 tax free from Alcoa each year which can be spent on eligible medical expenses (including but not limited to the plan’s cost sharing provisions). Unspent funds in the HRA can be rolled over the subsequent years, but any unspent balance is forfeited upon separation from Alcoa. The out-of-pocket maximums of all these options are also higher, but this is largely irrelevant empirically since it is extremely rare (less than 1%) for any employee to hit the out-of-pocket maximum of even the most generous option. Coinsurance rates are the same across all PPOs.

¹³Since, as noted, the prices of other coverage tiers are proportional to the family price, this is not an unreasonable

economically or statistically significant effect on the probability of choosing the outside good.

A.3. Extensions

In this section we briefly discuss several extensions to our application, following the discussion of possible extensions to our framework in Sections 2 and 3.

More than two coverage choices. As noted in Section 3, it would be conceptually straightforward to extend our empirical analysis to consider more than two choices. However, we face practical obstacles to doing so in our setting. In particular, as is typical in data sets like ours, we do not observe medical expenditures for employees covered by an HMO or who opted out of employer-provided coverage. We therefore cannot estimate the cost curve for these options. It is also difficult to model the demand for these two options, since the prices are not known, nor is it entirely clear how to define the “good” being purchased.¹⁴ We experimented with estimating demand and cost systems for the remaining five PPO options. However, the relatively small sample sizes on the other three PPO options combined with the relatively high multi-collinearity in relative prices among the different PPO options resulted in fairly imprecise (and therefore relatively uninformative) estimates of the demand and cost systems.

Moral hazard. As we discussed in Section 3.2, our framework also allows us to easily test for and quantify moral hazard, which is defined by the vertical distance between MC^H and MC^L . Moreover, as discussed in Section 2.4, when contract L provides partial coverage (as in our application) moral hazard will affect the welfare analysis. Therefore it is important to examine moral hazard empirically in our setting.

With two partial coverage contracts, c_i^H is defined as the incremental cost to the insurer of covering employee i with contract H rather than with contract L assuming i behaves as if he is covered by contract H . Analogously, c_i^L is the incremental cost to the insurer of covering employee i with contract H rather than with contract L assuming i behaves as if he is covered by contract L . Our foregoing estimates of AC , which were estimated on the sample of individuals who chose contract H , therefore gives us AC^H . And our estimate of MC , using our estimate of AC^H and our estimate of the demand curve for H (equation (11)), similarly gives us MC^H . To estimate AC^L we estimate the same cost equation (equation (12)) but on the sample of individuals who chose contract L . To back out MC^L from AC^L we use the demand curve for contract L , i.e. equation (11) estimated with D_i replaced by $1 - D_i$.

We have run this exercise on our baseline sample and were unable to reject the null of no moral hazard (i.e. $H_0 : MC^L = MC^H$). Our estimates were quite imprecise, suggesting that we may lack

approach.

¹⁴The price of the HMO is literally not known, and likely varies across geographic areas. Employees receive a \$1,000 “credit” if they opt out of any coverage. However, without knowing what price they face for purchasing insurance outside the company it is not clear what the true price is. Relatedly, in contrast to the PPO options, the characteristics of the HMO option and any coverage offered outside the firm are not known.

sufficient power in our setting to detect moral hazard. This may not be surprising given that the design of the insurance contracts in our setting (see Figure 3) should make moral hazard primarily affect those employees who expect to spend less than the contract L deductible. In practice, this is likely to be a small fraction of our data.¹⁵

As a different way to make this point, we applied the widely used moral hazard estimate of Manning et al. (1987)¹⁶ from the Rand Health Insurance Experiment to the total spending of each employee covered by contract H . We assumed a price effect which is based on the change in the marginal cost-sharing this employee would face under contract L compared to contract H , holding his realized (rather than expected) spending fixed. This back-of-the-envelope calculation led to an average change in insurer’s cost of 3%, driven by the fact that three quarters of the employees did not experience any change in marginal cost sharing. In light of this, we find it unsurprising that it is hard to detect moral hazard in this setting.

Departures from revealed preference. As we noted at the outset, our approach to welfare analysis has relied on revealed preferences. It is possible to use our framework for welfare analysis when we are not willing to assume revealed preferences, although this would require specification of the precise alternative choice model and how it maps to welfare. Some “behavioral” models are easily translated to our approach. Consider, for example, the possible role of defaults. The default option in our setting is contract L . If one believes that there is a (constant) fraction α of the sample who always chooses the default, then it is possible to implement our approach, and perform welfare analysis on the remaining $1 - \alpha$ share of the sample, who are “active” choosers.

¹⁵Considering in-network spending, there are 9% of the employees in our baseline sample who spend less than the contract L (in-network) deductible of \$500. Out-of-network spending would increase this share (but not by much).

¹⁶Manning, Willard G., Joseph P. Newhouse, Naihua Duan, Emmett B. Keeler, Arleen Leibowitz, and M. Susan Marquis, “Health insurance and the demand for medical care: Evidence from a randomized experiment,” *American Economic Review*, 77 (1987), 251-277.

Table A1: Summary statistics

	<u>2004 Company Data</u>					<u>March 2005 CPS</u>		
	All employees	Only salaried employees	Only salaried employees with new benefit design	Col. (3) limited to only employees who chose <i>H</i> or <i>L</i>	Col. (4) limited to employees with family coverage	All full-time employees	Only in manufacturing	White collar employees in manufacturing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Number of Individuals	36,814	11,964	11,325	7,263	3,779	83,118	11,178	4,688
Fraction Male	0.78	0.73	0.73	0.77	0.86	0.58	0.70	0.64
Fraction White	0.77	0.87	0.86	0.86	0.86	0.82	0.82	0.86
Fraction unionized	0.33	0.00	0.00	0.00	0.00	0.13	0.14	0.04
Age								
Mean	44.24	44.51	44.50	45.17	42.66	41.39	42.13	42.87
Std. Deviation	9.86	9.22	9.21	9.12	7.22	12.33	11.45	10.88
Median	45	45	45	46	43	41	42	43
Tenure with company (years)								
Mean	13.23	13.26	13.23	13.69	12.70	n/a	n/a	n/a
Std. Deviation	10.28	9.95	9.96	10.01	8.93	n/a	n/a	n/a
Median	11	12	12	13	12	n/a	n/a	n/a
Annual Salary (current \$US)								
Mean	53,103	71,622	72,821	74,017	80,999	41,869	46,195	63,157
Std. Deviation	47,642	77,936	79,373	91,530	112,790	47,955	45,435	58,072
Median	47,283	60,484	61,433	61,822	66,335	32,000	35,000	50,000

Columns (1) to (5) present summary statistics for different cuts of the 2004 Alcoa employees. Column (1) presents statistics for all active employees in our sample, column (2) for salaried employees only. Column (3) looks at a slightly smaller group of salaried employees who faced the new benefit design, and column (4) further restricts attention to salaried employees who chose either contract *H* or contract *L* (who are the primary focus of our analysis). Column (5) further limits the analysis to those who chose family coverage; this sample is used to generate our baseline estimates. For comparison, columns (6) to (8) present summary statistics for full time employees (defined as those who on average worked 35 or more hours per week in the previous year) in the March 2005 CPS. Column (6) shows all full time employees, column (7) shows all full time employees in manufacturing industries, and column (8) shows all full time white collar employees (defined based on occupation codes) in manufacturing industries; in these three columns we use CPS sampling weights (“earning weights” for the union variable, and “person weights” for all others).

Table A2: Robustness

Panel A: Welfare estimates from different specifications

	Competitive Equilibrium		Efficient Allocation		Welfare cost of Adverse Selection			
	Q	P	Q	P	In dollar terms (per market participant) ^a	Relative to social cost of efficient subsidy ^b	Relative to welfare cost of mandating contract H^c	Relative to total achievable welfare ^d
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1 Baseline (family coverage, no state fixed-effects)	0.617	463.51	0.756	263.94	9.55	21.1%	32.4%	3.4%
<i>Robustness to demand estimates</i>								
2 Probit demand	0.619	463.59	0.790	187.85	11.32	17.3%	31.6%	3.5%
3 Linear demand, constrained to go through (Q,P)=(1,\$0)	0.612	463.56	0.750	299.04	7.81	21.1%	30.2%	3.4%
4 Linear demand, constrained to go through (Q,P)=(0,\$800)	0.562	463.59	0.688	387.90	3.30	21.1%	16.4%	3.4%
5 Quadratic demand, constrained to go through (1,\$0) and (0,\$800)	0.587	463.58	0.738	343.51	5.00	18.8%	45.6%	4.5%
<i>Robustness to tax subsidy</i>								
6 Baseline specification, but accounting for pre-tax premiums	0.389	514.49	0.567	348.53	7.71	27.3%	16.8%	9.8%
<i>Robustness to sample and source of variation</i>								
7 State fixed-effects included (in both demand and cost regressions)	0.622	460.16	0.699	341.40	3.65	14.6%	6.5%	1.2%
8 State fixed-effects and demographics included (in both regressions)	0.641	440.00	0.724	306.67	4.42	15.3%	9.2%	1.3%
9 All coverage tiers, no state fixed-effects ^e	0.593	434.20	0.704	244.83	7.67	19.2%	14.2%	2.5%
10 Baseline specification, without the \$570 price group	0.641	460.57	0.740	202.09	9.77	17.0%	14.5%	1.8%

Panel B: Parameter estimates from different specifications

	Demand Equation				Average Cost Equation			
	alpha		beta		gamma		delta	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
1 Baseline (family coverage, no state fixed-effects)	0.940	(0.123)	-0.00070	(0.00032)	391.7	(26.8)	0.155	(0.064)
<i>Robustness to demand estimates</i>								
2 Probit demand	1.149	(0.316)	-0.00183	(0.00080)	----- same as Baseline -----			
3 Linear demand, constrained to go through (Q,P)=(1,\$0)	1.000	(imposed)	-0.00084	(0.00005)	----- same as Baseline -----			
4 Linear demand, constrained to go through (Q,P)=(0,\$800)	1.333	(imposed)	-0.00167	(0.00005)	----- same as Baseline -----			
5 Quadratic demand, constrained to go through (1,\$0) and (0,\$800) ^f	1.000	(imposed)	-0.00039	(imposed)	----- same as Baseline -----			
<i>Robustness to tax subsidy</i>								
6 Baseline specification, but accounting for pre-tax premiums	0.940	(0.123)	-0.00107	(0.00048)	391.7	(26.8)	0.239	(0.098)
<i>Robustness to sample and source of variation</i>								
7 State fixed-effects included (in both regressions)	0.919	(0.167)	-0.00065	(0.00040)	414.8	(37.0)	0.099	(0.090)
8 State fixed-effects and demographics included (in both regressions)	0.917	(0.170)	-0.00063	(0.00040)	394.5	(36.9)	0.104	(0.091)
9 All coverage tiers, no state fixed-effects ^e	0.848	(0.109)	-0.00059	(0.00032)	374.8	(22.8)	0.137	(0.062)
10 Baseline specification, without the \$570 price group	0.818	(0.195)	-0.00038	(0.00050)	406.3	(34.0)	0.118	(0.081)

Table reports results from alternative specifications. Panel B reports parameter estimates, and Panel A reports the (corresponding) implications for welfare analysis. Row 1 reports the baseline specification (previously described in the text), rows 2-5 report specifications that change the functional form of demand. Row 6 re-estimates the baseline specification with the price in both the demand and cost equation multiplied by 0.65 (one minus the average marginal tax rate in the sample). Row 7 includes state fixed effects in both the demand and cost equations, and row 8 also controls for employee characteristics (see Table 2). Row 9 increases the sample to include employees in all four coverage tiers. Row 10 tries to assess sensitivity to dropping the \$570 price group, which is the greatest outlier (see Figure 5). Standard errors (in parentheses) allow for an arbitrary variance-covariance matrix within each state.

^a Graphically, this is the area of triangle CDE (see Figure 1).

^b This is triangle CDE divided by $0.3Q_{eff}(P_{eq} - P_{eff})$.

^c Graphically, this is the area of triangle CDE divided by the area of triangle EGH (see Figure 1).

^d Graphically, this is the area of triangle CDE divided by the area of triangle ABE (see Figure 1).

^e N=7,263 for demand analysis, 4,622 for cost analysis; mean dependent variables are 0.64 (D) and \$424 (c), respectively. We include (de-meaned) indicator variables for the coverage tier in both the demand and cost equations (not shown); we multiply p and c by two for employees in the “employee only” coverage tier.

^f In the quadratic demand specification, the top reported coefficient of beta is the coefficient on the linear term, while the second is the coefficient on the quadratic term.

Table A3: Potential sample selection

Dependent variable:	1 if "outside good" was chosen, 0 otherwise		
	"Outside Good" does not include "opt out"		"Outside good" does include "opt out"
	Family coverage tier only (1)	All coverage tiers (2)	All coverage tiers (3)
Relative price	-0.000093 (0.00035) [0.98]	-0.000021 (0.00040) [0.96]	0.000002 (0.000003) [0.66]
Constant	0.287 (0.1580) [0.08]	0.292 (0.1150) [0.02]	0.296 (0.1580) [0.07]
Mean dependent variable	0.283	0.300	0.359
Number of obs.	5,271	10,386	11,325

The table reports results of estimating a variant of the demand equation shown in equation (11). The dependent variable is an indicator variable that takes the value of 1 if the employee chose any of the “outside options” and 0 if the employee chose either contract H or contract L . The “relative price” variable is, as in Table 2, the relative price of contract H compared to contract L . In columns (1) and (2) the “outside good” includes two lower coverage PPOs, a Health Reimbursement Account PPO, and an HMO. The sample in column (1) is limited to family coverage. The sample in column (2) includes all coverage tiers. We therefore include (de-meanned) indicator variables for the coverage tier (not shown) and multiply the price variable by two for employees in the “employee only” coverage tier. In column (3) the “outside good” definition is expanded to also include employees who opt out of coverage. Since coverage tier is not known for these employees, we include all employees regardless of coverage tier and do not include indicator variables for coverage tier. We define the price variable as the relative price associated with family coverage (regardless of the actual tier chosen, if known). All estimates are generated by OLS, standard errors (in parentheses) allow for an arbitrary variance-covariance matrix within each state, ; p-values are in [square brackets].