

Selection on Moral Hazard in Health Insurance[†]

By LIRAN EINAV, AMY FINKELSTEIN, STEPHEN P. RYAN,
PAUL SCHRIMPF, AND MARK R. CULLEN*

We use employee-level panel data from a single firm to explore the possibility that individuals may select insurance coverage in part based on their anticipated behavioral (“moral hazard”) response to insurance, a phenomenon we label “selection on moral hazard.” Using a model of plan choice and medical utilization, we present evidence of heterogenous moral hazard as well as selection on it, and explore some of its implications. For example, we show that, at least in our context, abstracting from selection on moral hazard could lead to overestimates of the spending reduction associated with introducing a high-deductible health insurance option. (JEL D82, G22, I13, J32)

Economic analysis of market failure in insurance markets tends to analyze selection and moral hazard as distinct phenomena. In this paper, we explore the potential for selection on moral hazard in insurance markets. By this we mean the possibility that moral hazard effects are heterogenous across individuals, and that individuals’ selection of insurance coverage is affected by their anticipated behavioral response to coverage. We examine these issues empirically in the context of employer-provided health insurance in the United States. Specifically, we break down the general

* Einav: Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072, and NBER (e-mail: leinav@stanford.edu); Finkelstein: Department of Economics, MIT, 50 Memorial Drive, Cambridge, MA 02142-1347, and NBER (e-mail: afink@mit.edu); Ryan: Department of Economics, University of Texas at Austin, 2225 Speedway Stop C3100, Austin, TX 78712-1690, and NBER (e-mail: sryan@atexas.edu); Schrimpf: Department of Economics, the University of British Columbia, 997-1873 East Mall, Vancouver, BC, Canada V6T 1Z1 (e-mail: schrimpf@mail.ubc.ca); Cullen: Department of Internal Medicine, School of Medicine, Stanford University, 1265 Welch Road, Stanford, CA 94305-5414, and NBER (e-mail: mrcullen@stanford.edu). 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, Michael Powell, Iuliana Pascu, and James Wang for outstanding research assistance; and to Ben Handel, Justine Hastings, Jim Heckman, Igal Hendel, Nathan Hendren, Kate Ho, Pat Kline, Jon Levin, Matt Notowidigdo, Phil Reny, Rob Townsend, numerous seminar participants, and four anonymous referees for helpful comments and suggestions. 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), the John D. and Catherine T. MacArthur Foundation Network on Socioeconomic Status and Health, and Alcoa, Inc. (Cullen), and the US Social Security Administration through grant #5 RRC08098400-03-00 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. Einav also acknowledges the hospitality of the Center for Advanced Study in the Behavioral Sciences at Stanford. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

[†] To view additional materials, visit the article page at <http://dx.doi.org/10.1257/aer.103.1.178>.

problem of adverse selection to two components: one is driven by the “traditional” selection on the *level* of expected health risk, while the other is driven by *slope* of spending, namely the incremental medical utilization that is due to greater insurance coverage, which we refer to as “moral hazard.”

Such selection on moral hazard has implications for the standard analysis of both selection and moral hazard. For example, a standard approach to influence selection in insurance markets is risk adjustment; i.e., pricing on observable characteristics that predict one’s insurance claims. The potential for selection on moral hazard, however, suggests that monitoring techniques that are usually thought of as reducing moral hazard—such as cost sharing that varies across categories of claims with differential scope for behavioral response—may also have important benefits in combatting adverse selection. In contrast, a standard approach to influencing moral hazard is to offer plans with higher consumer cost sharing. But if individuals’ anticipated behavioral response to coverage affects their propensity to select such plans, the magnitude of the behavioral response could be much lower (or much higher) from what would be achieved if plan choices were unrelated to the behavioral response. As we discuss in more detail below, not only the existence of selection on moral hazard but also the sign of any relationship between anticipated behavioral response and demand for higher coverage is *ex ante* ambiguous. Ultimately, these are open empirical questions.

Health insurance provides a particularly interesting setting in which to explore these issues. Both selection and moral hazard have been well documented in the context of employer-provided health insurance. Moreover, given the extensive government involvement in health insurance, as well as the concern about the size and rapid growth of the health care sector, there is considerable academic and public policy interest in a better understanding of selection and moral hazard in this context.

Recognition of the possibility of selection on moral hazard, however, highlights potentially important limitations of analyzing these problems in isolation. For example, the sizable empirical literature on the likely spending reductions that could be achieved through higher consumer cost sharing has intentionally focused on isolating and exploring exogenous changes in cost sharing, such as those induced by the famous RAND experiment (Manning et al. 1987; Newhouse 1993). Yet, the very same feature that solves the causal inference problem—namely, randomization (or attempts to approximate it in the subsequent quasi-experimental literature on this topic)—removes the endogenous choice element. It thus abstracts, by design, from any selection on moral hazard, which could have important implications for the spending reductions achieved through offering plans with higher consumer cost sharing. This is particularly relevant since substantial plan choice is now the norm not only in private health insurance but also, increasingly, in public health insurance programs, such as Medicare Part D.

We begin by presenting a utility-maximizing model of individual health insurance plan choices and subsequent health care spending. The model characterizes individuals as associated with two distinct risk attributes: a “level” and a “slope.” The former refers to their health risk, or their expected level of health care spending without insurance. The latter captures the incremental health care spending from insurance coverage (i.e., the slope of health care spending with respect to its out-of-pocket price). The use of the term “moral hazard” is far from standard in the

literature, and the model allows us to be precise as to what we mean by it. We define moral hazard as the slope of health care spending (with respect to price), and by “selection on moral hazard” we refer to the component of adverse selection that is driven by heterogeneity in this slope parameter. In other words, while traditional models of adverse selection focus on heterogeneity in (and selection on) the level of expected medical utilization, we emphasize that adverse selection could also be driven by selection on the slope of medical spending with respect to its price. That is, greater coverage would be more attractive for individuals whose health care utilization would increase more sharply in response to this coverage, thus generating greater cost to the insurance company.

We explore these issues empirically in the specific setting of the US employees at Alcoa, Inc., a large multinational producer of aluminum and related products. Naturally, as we emphasize below, this makes our quantitative results specific to our setting. We have individual-level panel data on the health insurance options, choices, and subsequent medical utilization of Alcoa’s employees (and their covered dependents). Crucially for identifying and estimating moral hazard, we observe variation in the health insurance options offered to different groups of workers at different points of time stemming from staggered timing of new union contracts. We present descriptive and motivating difference-in-differences estimates on moral hazard in our setting, as well as patterns that may be suggestive of heterogeneity in and selection on this moral hazard effect.

In order to formalize the analysis of selection on moral hazard and to explore some of its implications, we embed the economic model of coverage choice and health care spending within an econometric model that allows for unobserved heterogeneity across individuals along three dimensions—health expectations, risk aversion, and moral hazard—and for flexible correlation across these three. All else equal, willingness to pay for coverage is increasing in the individual’s health expectation and his risk aversion; these are standard results. In addition, in our model, willingness to pay for coverage is increasing in the individual’s moral hazard: individuals with a greater behavioral response to coverage benefit more from greater coverage, since they will consume more care as a result. This is the selection on moral hazard comparative static that is the focus of our paper. Empirically, however, the sign (let alone the magnitude) of any selection on moral hazard is ambiguous and depends on the extent of heterogeneity in moral hazard as well as the correlation between moral hazard and the other primitives that affect health insurance choice, expected health, and risk aversion. We use this model and the data to recover the joint distribution of individuals’ (unobserved) health, risk aversion, and moral hazard. The model is estimated using Markov Chain Monte Carlo Gibbs sampler, and its fit appears reasonable.

Qualitatively, we find that individuals who exhibit a greater behavioral response to coverage are more likely to choose higher coverage plans. Quantitatively, we estimate substantial heterogeneity in moral hazard and selection on it. We focus on the counterfactual of moving from the most comprehensive to the least comprehensive of the new options—essentially moving individuals from a no-deductible plan to a high (\$3,000 for family coverage) deductible plan. In terms of heterogeneity in moral hazard, we find that the standard deviation across individuals of the spending reduction from this change in plans is more than twice the average. In terms of

selection on moral hazard, we find that for determining the choice between these two plans, selection on moral hazard is roughly as important as selection on health risk, and considerably more important than selection on risk aversion.

We use the model to examine some of the implications of the selection on moral hazard we detect. For example, our results suggest that the spending reduction associated with introducing a high-deductible plan could be substantially lower than what would be predicted if we were to ignore selection on moral hazard and assume that those who choose the high-deductible plan are drawn randomly. This is a direct consequence of our finding that those who select less comprehensive coverage are likely to exhibit a smaller behavioral response to the insurance coverage.

Our paper is related to several distinct literatures. Our modeling approach is closely related to that of Cardon and Hendel (2001), which is also the approach taken by Bajari et al. (2010), Carlin and Town (2010), and Handel (2012) in modeling health insurance plan choice. Like us, all of these papers have allowed for selection based on expected health risk. Our paper differs in our focus on identifying and estimating moral hazard—and in particular, heterogeneous moral hazard—and in examining the relationship between moral hazard and plan choice. From a methodological perspective, we also differ from these and many other discrete choice models in that we do not allow for a choice-specific, i.i.d. error term, which does not seem appealing given the vertically rankable nature of our choices.

Our difference-in-differences analysis of the spending reduction associated with changes in cost sharing is related to a sizable experimental and quasi-experimental literature in health economics analyzing the impact of higher consumer cost sharing on spending. (See Chandra, Gruber, and McKnight 2010 for a recent review.) Our subsequent exploration of heterogeneity in this average moral hazard effect and selection on it, however, suggests the need for caution in using such estimates, which do not account for endogenous plan selection when forecasting the likely spending effects of introducing the option of plans with higher consumer cost sharing. It also suggests that one can embed the basic identification approach of the difference-in-differences framework in a model that allows for and investigates such endogenous selection.

Our examination of selection on moral hazard is motivated in part by the growing empirical literature demonstrating that focusing on selection that is driven by one-dimensional heterogeneity in risk, as in the early seminal theoretical contributions to the topic, may miss many interesting aspects of actual markets. This literature has tended to abstract from moral hazard, and focused on selection on preferences, such as risk aversion (Finkelstein and McGarry 2006; Cohen and Einav 2007), cognition (Fang, Keane, and Silverman 2008), or desire for wealth after death (Einav, Finkelstein, and Schrimpf 2010). Our exploration highlights another potential dimension of selection, and one that has particularly interesting implications for contract design (in contexts where moral hazard is important). While there are questions for which the extent to which selection occurs on the basis of expected health or moral hazard does not matter (see, e.g., Einav, Finkelstein, and Cullen 2010), we illustrate in this paper that breaking down selection to “selection on levels” and “selection on slopes” can be important for answering questions regarding the design of contracts to reduce selection and the implications of contract design for spending.

Despite its potential importance, we are not aware of any empirical work attempting to identify and analyze selection on moral hazard in insurance markets. The basic idea of selection on moral hazard, however, is not unique to us. Similar ideas have appeared in several other contexts.¹ Indeed, one general way to think about the concept of selection on moral hazard is in the context of estimating a treatment effect of insurance coverage on medical expenditure. As we already hinted, within such a framework, selection on health risk would be equivalent to heterogeneity in (and selection on) the *level* (or constant term in a regression of medical spending on insurance coverage), while selection on moral hazard can be thought of as heterogeneity in (and selection on) the *slope* coefficient. Heckman, Urzua, and Vytlačil (2006) present an econometric examination of the properties of IV estimators when individuals select into treatment in part based on their anticipated response to the treatment, a phenomenon they refer to as “essential heterogeneity.” They subsequently apply these ideas in the context of the returns to education in Carneiro, Heckman, and Vytlačil (2011).

The rest of the paper proceeds as follows. Section I develops our model of an individual’s health insurance plan choice and spending decisions. Section II describes the data and Section III presents descriptive evidence of moral hazard. In Section IV we present the econometric specification of our model and describe its identification and estimation. Section V presents our main results, including some of their implications for spending. Section VI presents some illustrative welfare analysis based on the estimates. The last section concludes.

I. A Model

We begin by presenting a stylized model of individual coverage choice and health care utilization. The model allows us to more precisely define the objects that we focus on, “moral hazard” and “selection on moral hazard.” The model will also be the main ingredient in our subsequent econometric specification and counterfactual exercises.

A. A Model of Coverage Choice and Utilization

We consider a two-period model that is designed to allow us to isolate and examine three potential determinants of insurance coverage choice: risk aversion, expected health care needs, and the incremental health spending that is driven by insurance coverage. In the first period, a risk-averse expected utility–maximizing individual makes an optimal health insurance coverage choice, using his available information to form his expectation regarding his subsequent health realization. In the second period, the individual observes his realized health and makes an optimal health care utilization decision, which depends on realized health as well as his coverage.²

¹For example, in the context of appliance choices and phone plan choices, respectively, Dubin and McFadden (1984) and Miravete (2003) estimate models in which the choice is allowed to depend on subsequent utilization, which in turn may respond to the utilization price.

²Existing work in this area (Cardon and Hendel 2001; Bajari et al. 2010; Carlin and Town 2010; and Handel 2012) followed a similar modeling approach. While tractable, the two-period model abstracts from the fact that utilization decisions are made throughout the coverage year, before the uncertainty about subsequent health is fully

We begin with notation. This is a model of individual behavior, so we omit i subscripts to simplify notation; in Section IV we describe how individuals may vary. At the time of his utilization choice (period 2), an individual is characterized by two objects: his health realization λ , and his price sensitivity ω . λ captures the uncertain aspect of demand for health care, with higher λ representing sicker individuals who demand greater health care consumption. The parameter ω determines how responsive health care utilization decisions are to insurance coverage. In other words, ω affects the individual's price elasticity of demand for health care with respect to its (out-of-pocket) price; individuals with higher ω increase their utilization more sharply in response to more generous insurance coverage. The focus of the paper will be on how plan choice varies with ω . That is, ω is the object that we refer to as moral hazard, although we defer a more detailed discussion of its interpretation until after the description of the model.

At the time of coverage choice (period 1), an individual is characterized by three objects: $F_\lambda(\cdot)$, ω , and ψ . The first, $F_\lambda(\cdot)$, represents the individual's expectation about his subsequent health risk λ . It is precisely because individuals do not know λ with certainty at the time of coverage choice that they demand insurance. The second object is ω , which determines the individual's period 2 price elasticity of demand for health care. Because individuals are forward-looking, they anticipate that ω will subsequently affect their utilization choices, and this in turn affects their utility from different coverages. It is this channel that creates the potential for selection on moral hazard, which is the main focus of our paper. Finally, the third object is ψ , which captures the individual's coefficient of absolute risk aversion. Importantly, unlike ω and $F_\lambda(\cdot)$, which enter the coverage choice and also affect utilization decisions, risk preferences affect coverage choice but do not directly affect utilization.

Utilization Choice.—In the second period, insurance coverage, denoted by j , is taken as given. We assume that the individual's health care utilization decision is made in order to maximize a trade-off between health and money. Specifically, we assume that the individual's second-period utility is separable in health and money and can be written as $u(m; \lambda, \omega) = h(m - \lambda; \omega) + y(m)$, where $m \geq 0$ is the monetized utilization choice, λ is the monetized health realization, and $y(m)$ is the residual income. Naturally, $y(m)$ is decreasing in m at a rate that depends on coverage. We further assume that $h(m - \lambda; \omega)$ is concave in its first argument: it is increasing for low levels of utilization (when treatment improves health) and is decreasing eventually (when there is no further health benefit from treatment and time costs dominate). Thus, the marginal benefit from incremental utilization is decreasing. Using this formulation, we think of λ , the underlying health realization, as shifting the level of optimal utilization m^* . Finally, we assume that $h(m - \lambda; \omega)$ is increasing in its second argument, but this is purely a normalization that (as we will see below) allows us to interpret individuals with higher ω as those who are more elastic with respect to the price of medical utilization.

resolved. In more recent work, we explore in more detail the implications of this latter aspect of health insurance (Aron-Dine et al. 2012).

We parametrize further so that the second-period utility function is given by

$$(1) \quad u(m; \lambda, \omega, j) = \underbrace{\left[(m - \lambda) - \frac{1}{2\omega}(m - \lambda)^2 \right]}_{h(m - \lambda; \omega)} + \underbrace{[y - c_j(m) - p_j]}_{y(m)}.$$

That is, we assume that $h(m - \lambda; \omega)$ is quadratic in its first argument, with ω affecting its curvature. We also explicitly write the residual income as the initial income y minus the premium p_j associated with coverage j and the out-of-pocket expenditure $c_j(m)$ associated with utilization m under coverage j . Given this parameterization, the optimal utilization is given by

$$(2) \quad m^*(\lambda, \omega, j) = \arg \max_{m \geq 0} u(m; \lambda, \omega, j),$$

which gives rise to the realized utility $u^*(\lambda, \omega, j) \equiv u(m^*(\lambda, \omega, j); \lambda, \omega, j)$.

To facilitate intuition, we consider here optimal utilization for the case of a linear coverage contract, so that $c_j(m) = c \cdot m$, with $c \in [0, 1]$ (in the empirical application below we explicitly account for the nonlinear coverage contracts offered in the data). Full insurance is given by $c = 0$ and no insurance is given by $c = 1$. The first-order condition implied by equation (2) is given by $1 - (m - \lambda)/\omega - c = 0$, or

$$(3) \quad m^*(\lambda, \omega, c) = \max[0, \lambda + \omega(1 - c)].$$

Abstracting from the potential truncation of utilization at zero, the individual will optimally choose $m^* = \lambda$ with no insurance (when $c = 1$) and $m^* = \lambda + \omega$ with full insurance (when $c = 0$). Thus, ω can be thought of as the incremental utilization that is attributed to the change in coverage from no insurance to full insurance. One way to think about this model, therefore, is that λ represents nondiscretionary health care shocks that the individual will pay to treat, regardless of insurance. In addition, there exist discretionary health care utilizations that will not be undertaken without insurance. With insurance, some amount of this discretionary care will be consumed, and this incremental amount is increasing in ω .³

Coverage Choice.—In the first period, the individual faces a fairly standard insurance coverage choice. As mentioned, we assume that the individual is an expected-utility maximizer, with a coefficient of absolute risk aversion of ψ . We further assume that the individual's von Neumann Morgenstern (vNM) utility function is of the constant absolute risk aversion (CARA) form $w(x) = -\exp(-\psi x)$. In a typical insurance setting, $w(x)$ is defined solely over financial outcomes. Because moral hazard is present, however, $w(x)$ is defined over the realized second-period utility $u^*(\lambda, \omega, j)$.

³We have written the model as if it is the individual who makes all the utilization decisions. To the extent that physicians also respond to the individual's coverage (and they are likely to), our interpretation of the utilization choice should be thought of as some combination of both the individual's and the physician's decisions.

Consider now a set of coverage options J , with each option $j \in J$ defined by its premium p_j and coverage function $c_j(m)$. Following the above assumptions, the individual will then evaluate his expected utility from each option,

$$(4) \quad v_j(F_\lambda(\cdot), \omega, \psi) = - \int \exp(-\psi u^*(\lambda, \omega, j)) dF_\lambda(\lambda),$$

with his optimal coverage choice given by

$$(5) \quad j^*(F_\lambda(\cdot), \omega, \psi) = \arg \max_{j \in J} v_j(F_\lambda(\cdot), \omega, \psi).$$

Willingness to pay for more coverage is generally increasing in risk aversion ψ and in risk $F_\lambda(\cdot)$ (in a first-order stochastic dominance sense).⁴ Given our specific parametrization, willingness to pay for more coverage is also increasing in ω , thus possibly generating what we term as selection on moral hazard.⁵

B. Interpreting ω as Moral Hazard

As noted, our focus is on the parameter ω and on its importance in driving coverage choice. We think of ω as moral hazard and of its relationship with coverage choice as selection on moral hazard. Given the varying ways by which the term “moral hazard” has been used (and abused) in economics in general and in the context of health insurance in particular, it seems useful to discuss the interpretation of ω and why we think it may be appropriate to refer to it as moral hazard.

Traditional models of adverse selection in health insurance focus on the possibility that sicker individuals will choose greater health insurance coverage. This source of selection is captured in our model by the fact that individuals with greater $F_\lambda(\cdot)$ (in a first-order stochastic dominance sense) purchase greater coverage. The key conceptual distinction we are interested in is the possibility that selection may be driven not only by the level of expected health care spending, but also by its “slope,” or by how health care spending changes with insurance coverage. In other words, we are interested in health insurance choices (selection) that are affected by the incremental medical expenditure associated with increased coverage. We refer to this incremental spending, captured in the model by ω , as moral hazard and to selection on it as selection on moral hazard. Just like traditional selection, which would lead to adverse selection (sicker individuals are willing to pay more for insurance, and at the same time are associated with greater expected cost to the insurance company),

⁴These comparative statics do not always hold. The model has unappealing properties when a significant portion of the distribution of λ is over the negative range, in which case the individual is exposed to a somewhat artificial uninsurable (background) risk (since spending is truncated at zero). We are not particularly concerned about this feature, however, as our estimated parameters do not give rise to it, and because we have experimented with a (non-elegant) modification to the model that does not have this feature, and the overall results were similar.

⁵In a more general model, ω is associated with two effects. One is the increased utilization, which increases willingness to pay. The second effect is the increased flexibility to adjust utilization as a function of the realized uncertainty (λ), which in turn reduces risk exposure and reduces willingness to pay for insurance. Our specific parameterization was designed to have spending under no insurance unaffected by ω ; this eliminates the latter effect, and therefore makes the comparative statics unambiguous.

in our model selection on moral hazard is also adverse in the sense that higher-moral hazard individuals are willing to pay more for the same amount of coverage and will also be more expensive for the insurance company. Thus, we view selection on moral hazard as one possible component of the overall adverse selection.

The use of the term “moral hazard” to refer to the responsiveness of health care spending to insurance coverage dates back at least to Arrow (1963). Consistent with the notion of hidden action, which is typically associated with the term “moral hazard,” it has been conjectured that health insurance may induce individuals to exert less (unobserved) effort in maintaining their health (e.g., Ehrlich and Becker 1972). In the context of health insurance, however, moral hazard is often used to refer to the price elasticity of demand for health care, conditional on underlying health status (Pauly 1968; Cutler and Zeckhauser 2000). We thus follow this abuse of terminology and use the term in a similar way. In other words, our model, like most in this literature, does not consider the potential impact of insurance on underlying health λ . As a result, the asymmetric information problem that we associate with moral hazard is arguably more accurately described as one of hidden information (rather than of hidden action). The individual’s actions (utilization) are observed and contractible, but his underlying health λ and his price sensitivity ω are unknown to the insurer. For our purposes, whether the problem is one of hidden information or hidden action is simply an issue of appropriate usage of terminology, and here we simply follow convention.

Our model is designed for conceptual clarity and analytical tractability, both of which come at the cost of not explicitly modeling the underlying primitives that give rise to ω . An individual’s incremental utilization response to increased insurance coverage (ω) is presumably driven by a number of “deeper” primitives including his value of time (income) and disutility of doctor visits. It may also relate to the severity and nature of his underlying health conditions—some of which are more likely to be price inelastic than others—as well as to one’s risk aversion regarding future health conditions.⁶

Also for clarity and tractability, we chose to model ω as a level shift in spending that is (except due to the truncation of spending at zero) independent of one’s health (λ) (see, e.g., equation (3)). Our choice of the utility function in equation (1) is designed to achieve a simple economic interpretation of the key parameters of interest in the first-order condition (3), so that λ (health status) is the monetized health spending without insurance (i.e., one’s nondiscretionary spending) and ω captures incremental, discretionary spending as individuals are moved from no insurance to full insurance. This allows us to straightforwardly measure and compare the magnitude of (and heterogeneity in) health risk λ and moral hazard ω .

It is not a priori obvious whether or not moral hazard affects individuals in a manner that is additively separable from their health. On the one hand, it seems plausible that seeking care for a minor skin irritation may be unaffected by one’s overall severity of illness. On the other hand, one could also imagine that changes

⁶We have modeled the second period utility in a static way, with no uncertainty. As a result, moral hazard is not directly determined by risk aversion. Nonetheless, one can imagine that more risk-averse people might be less sensitive to price in making their medical care consumption decisions, making them have a lower ω in the context of our model. Our empirical specification below will therefore allow for an arbitrary correlation between ω and risk aversion (ψ).

in medical care utilization in response to insurance coverage would depend on one's underlying health; for example, sicker individuals have more occasions to exercise moral hazard.⁷ In principle, our setting does not preclude this. Although we do not explicitly model this complementarity "within" an individual, our empirical specification below will allow for this relationship by modeling a cross-sectional distribution with an arbitrary correlation between moral hazard and health risk. Thus, a multiplicative model, for example, can be approximated by simply relabeling a multiplicative moral hazard effect ω' to be equal to ω/λ . The key modeling assumption is therefore not the additive separable relationship, but rather the fact that all uncertainty at the time of coverage choice is about health (λ), while moral hazard (ω) is assumed to be known at the time of coverage choice.⁸

II. Setting and Data

Baseline Sample.—We study health insurance choices and medical care utilization of the United States-based workers (and their dependents) at Alcoa, Inc., a large multinational producer of aluminum and related products. Our main analysis is based on data from 2003 and 2004, although for some of the analyses we extend the sample through 2006.

Our data contain the menu of health insurance options available to each employee, the employee's coverage choices, and detailed, claim-level information on his (and any covered dependents') medical care utilization and expenditures for the year.⁹ The data also contain demographic information, including the employee's union affiliation, employment type (hourly or salary), age, race, gender, annual earnings, job tenure at the company, and the number and ages of other insured family members. In addition, we obtained a summary proxy of an individual's health based on software that predicts future medical spending on the basis of previous years' detailed medical diagnoses and claims, as well as basic demographics (age and gender); importantly for our purposes, this generated health risk score is not a function of the individual's coverage choice.¹⁰

In 2004, Alcoa introduced a new set of health insurance Preferred Provider Organization (PPO) options in an effort to control health care spending by encouraging employees to move into plans with substantially higher consumer cost sharing. The new options were introduced gradually to different employees based on

⁷The findings in our data are consistent with a model in which moral hazard effects are not multiplicative in underlying health. Specifically, we find that changes in health care coverage are associated with changes in doctor and outpatient utilization but not with (the much more expensive) inpatient utilization. (See online Appendix B.)

⁸Alternative models could make moral hazard stochastic at the time of coverage choice, but would come at the cost of either equally strong assumptions or reliance on functional form for identification. In online Appendix E we report results from one such model.

⁹Health insurance choices are made in November, during the open enrollment period, and apply for the subsequent calendar year. They can be changed during the year only if the employee has a qualifying event, which is not common.

¹⁰This is a relatively sophisticated way of predicting medical spending as it takes into account the differential persistence of different types of medical claims (e.g., diabetes versus car accident) in addition to overall utilization, demographics, and a rich set of interactions among these measures. The particular software we use is a risk adjustment tool called DXCG Risk Solution that was developed by Verisk Health and is used by, among other organizations, the Center for Medicare and Medicaid services in determining reimbursement rates in Medicare Advantage. See Bundorf, Levin, and Mahoney (forthcoming), Carlin and Town (2010), and Handel (2012) for other examples of academic uses of this type of predictive diagnostic software.

TABLE 1—SUMMARY STATISTICS FOR 2003 SAMPLE

	Baseline sample (1)	Switched			
		In 2004 (2)	In 2005 (3)	In 2006 (4)	After 2006 (5)
Observations	3,995	682	974	1,075	1,264
Average age	41.3	44.5	39.7	38.3	43.3
Average annual income	31,292	39,715	25,532	29,952	32,316
Average tenure with Alcoa	10.2	15.5	8.2	5.7	12.7
Fraction male	0.84	0.96	0.73	0.86	0.85
Fraction white	0.72	0.85	0.44	0.82	0.79
Fraction single coverage	0.23	0.21	0.25	0.23	0.22
Average health risk score ^a	0.95	1.06	0.91	0.86	1.01
Avg. number of insured family members (if > 1)	2.8	2.7	2.8	2.9	2.6
Total annual medical spending (US\$) ^b	5,283	5,194	5,364	5,927	4,717

Notes: Column 1 presents statistics based on the 2003 data for our baseline sample, which covers all hourly union workers not covered by the Master Steelworker's Agreement (except those that get dropped in the process of the data cleaning described in the text). Columns 2–5 partition our baseline sample based on the year in which employees were switched to the new set of health insurance options.

^aHealth risk score is normalized, so that 1 indicates the expected medical expenditure for a random draw from a nationally representative under-65 population. To construct this table, we assign each employee the average risk score of all covered family members.

^bTotal annual medical spending is for employees and any covered dependents.

their union affiliation, since new benefits could only be introduced when an existing union contract expired. The staggered timing in the transition from one set of insurance options to another provides a plausibly exogenous source of variation that can help us identify the impact of health insurance on medical care utilization. To use this variation, we focus attention on the approximately 4,000 unionized workers (each year), who belong to one of 28 different unions whose benefit could only be introduced at contract expiration. Online Appendix A provides additional details on the construction of this baseline sample.

Column 1 of Table 1 provides some summary statistics of our baseline sample in 2003. Our sample is 72 percent white, 84 percent male, with an average age of 41, average annual income of about \$31,000, and an average tenure of about 10 years at the company. Approximately one-quarter of the sample has single (employee-only) coverage, while the rest cover additional dependents. The health risk score is calibrated to be interpreted as predicted medical spending relative to a randomly drawn person under 65 in the nationally representative population; on average, individuals in our sample have predicted medical spending that is about 5 percent lower than this benchmark. The remaining columns of Table 1 show summary statistics for four different groups of employees based on when they were switched to the new benefit options; we discuss this comparison when we present our difference-in-differences strategy and results below.

As noted, our main analysis is based on the 2003 and 2004 data (7,570 employee-years and 4,477 unique employees). We exclude the 2005 and 2006 data from our primary analysis because it introduces two challenges for estimation of our plan choice model. First, the relative price of comprehensive coverage on the new options was raised substantially in 2005 and raised further in 2006, yet remarkably few employees already in the new option set changed their plans. Second, the pricing in 2006 makes some of the observed choices clearly dominated. Both these patterns




FIGURE 1. THE CROSS-SECTIONAL DISTRIBUTION OF MEDICAL EXPENDITURE

Notes: The figure presents the distribution of total annual medical expenditure for each employee (and any covered dependents) in our baseline sample. The graph uses a log scale, such that the second bin covers expenditure lower than $\exp(0.5)$, the next covers expenditures between $\exp(0.5)$ and $\exp(1)$, and so on; the x -axis labels show the corresponding dollar amounts of selected bins. An observation is an employee-year, pooling data from 2003 and 2004. The gray bars correspond to employees with a single coverage, while the black bars correspond to employees who also covered additional dependents (spouse, children, or both).

are consistent with substantial evidence of inertial behavior in health insurance plan choices (Carlin and Town 2010; Handel 2012). Rather than modeling this behavior, we prefer to restrict the data to a time period where it is less central to understanding plan choices.

The main drawback to limiting the data to 2003 and 2004 is that less than one-fifth of our sample were offered the new benefits starting in 2004, while another half of the sample was transitioned to the new benefits in 2005 and 2006 (Table 1, top row). Therefore, for some of the descriptive evidence (which does not require an explicit model of plan choice), we use data from 2003–2006, which produces qualitatively similar descriptive results but with greater precision.

Medical Spending.—We have detailed, claim-level information on medical expenditures and utilization. Our primary use of these data is to construct annual total medical spending for each employee (and his covered dependents); in online Appendix B, we also use these data in a less aggregated way to break out spending by category (doctor’s office, outpatient, inpatient, and other). Figure 1 graphs the distribution of medical spending for our sample. We show the distribution separately for the approximately three-quarters of our sample with nonsingle coverage and the remainder with single-employee coverage. Not surprisingly, average spending is substantially higher in the former group. Across all employees, the average annual

spending (on themselves and their covered dependents) is about \$5,200.¹¹ As is typical, medical expenditures are extremely skewed. For example, for nonsingle coverage, average spending (\$6,100) is about 2.5 times greater than the median spending (\$2,400), and about 4 percent of our baseline sample has no spending, while each of the employees in the top decile spends over \$13,000.

Health Insurance Options and Choices.—A very attractive feature of our setting is that the PPO plans we study differ—across the new and old regimes and within each regime—only in their consumer cost sharing requirements. They are identical on all non-cost sharing features, such as the network definition. Table 2 summarizes the original and new plan options and the fraction of employees who choose each option in our baseline sample. Employees may choose from up to four coverage tiers: single (employee-only) coverage, or one of three nonsingle coverage tiers (employee plus spouse, employee plus children, or family). In our analysis, we take coverage tier as given, assuming that it is driven primarily by family structure.¹²

There were three PPO options under the old benefits and five entirely different PPO options under the new benefits. Because there was no option of “staying in your existing plan”—the five new options were all distinct from the three old options in both their names and their designs—individuals did not have the option of passively being defaulted into their existing coverage. We show in Table 3 that plan choices for those who are switched to the new options are also consistent with the notion of “active” choices. As a result, we suspect that defaults did not play an important role in the choice of new benefits. Indeed, although option 4 was the default coverage option, it was not the most common choice (Table 2).

The primary change from the old to the new benefits was to offer plans with higher deductibles and to increase the lowest out-of-pocket maximum.¹³ As shown in the table, under the new options there was a shift to plans with higher consumer cost sharing. Under the old options, virtually all employees faced no deductible. Looking at employees with nonsingle coverage in panel B (patterns for single-coverage employees are similar), about two-fifths faced a \$2,000 out-of-pocket maximum while three-fifths faced a \$5,000 out-of-pocket maximum. By contrast, under the new options, about a third of the employees faced a deductible, and all of them faced a high out-of-pocket maximum of at least \$5,000 for nonsingle coverage.¹⁴

¹¹ A little over one-quarter of total spending is in doctor offices, about one-third is for inpatient hospitalizations, and about one-third is for outpatient services. About half of the remaining 4 percent of spending is accounted for by emergency room visits.

¹² Employee premiums vary across the four coverage tiers according to fixed ratios. Cost-sharing provisions differ only between single and nonsingle coverage. Specifically, for a given PPO, deductibles and out-of-pocket maxima are twice as great for any nonsingle coverage tier as they are for single coverage. As shown in Table 1, about one-quarter of the sample chooses single coverage. Within nonsingle coverage, slightly over half choose family coverage, 30 percent choose employee plus spouse, and about 16 percent choose employee plus children (not shown).

¹³ At a point in time, prices within a coverage tier vary slightly across employees (in the range of several hundred dollars) under either the old or new options, depending on the employee's affiliation (see Einav, Finkelstein, and Cullen 2010 for more detail). Premiums were constant over time under the old options; as mentioned, under the new options, premiums were increased substantially (and cross-employee differences were removed) in 2005 and 2006 (not shown).

¹⁴ A \$5,000 (\$2,500) out-of-pocket maximum for nonsingle (single) coverage is rarely binding. With no deductible and a 10 percent consumer cost sharing, the employee must have \$50,000 (\$25,000) in total annual medical expenditures to hit this out-of-pocket maximum. Using the realized claims, we calculate that only about 1 percent of the employees would hit the out-of-pocket maximum in a given year. By contrast, under the old options the

TABLE 2—OLD AND NEW HEALTH PLANS

	Original plan options			New plan options				
	Option 1	Option 2	Option 3	Option 1 ^a	Option 2	Option 3	Option 4	Option 5
<i>Panel A. Single coverage (N = 1,679)</i>								
Plan features:								
Deductible	1,000	0	0	1,500	750	500	250	0
Out-of-pocket maximum	5,000	2,500	1,000	4,500	3,750	3,500	2,750	2,500
Average share of spending paid out of pocket ^b	0.580	0.150	0.111	0.819	0.724	0.660	0.535	0.112
Employee premium ^c	0	351	1,222	0	132	224	336	496
Fraction choosing each option ^d	3.3%	63.5%	33.2%	14.1%	0.0%	2.2%	37.8%	45.9%
<i>Panel B. Nonsingle coverage (N = 5,895)</i>								
Plan features:								
Deductible	2,000	0	0	3,000	1,500	1,000	500	0
Out-of-pocket maximum	10,000	5,000	2,000	9,000	7,500	7,000	5,500	5,000
Average share of spending paid out of pocket ^b	0.495	0.130	0.098	0.732	0.600	0.520	0.387	0.111
Employee premium ^c	0	354	1,297	0	364	620	914	1,306
Fraction choosing each option ^d	0.6%	56.1%	43.3%	3.9%	0.6%	1.8%	24.4%	69.3%

Notes: The table summarizes the key features of the original and new health insurance coverage options. The features shown apply to in-network spending. Not shown are coinsurance rates (applied to those who reached the deductible but have yet to reach the out-of-pocket maximum), which are 10 percent in all plans (old and new). There are some other small differences between the original and new options that are associated with out-of-network spending, preventive care, and certain treatments associated with copays rather than coinsurance in the original set of options. See text for further details.

^a The New Option 1 includes a health reimbursement arrangement (HRA). Every year the employer sets aside \$750 (for single; \$1,250 for nonsingle coverage) that the employee can use (tax free) to pay for a variety of expenses such as deductibles and coinsurance payments. Unused HRA funds roll over to future years and, eventually, can be used during retirement to finance health insurance, provided through the company or through COBRA. Our baseline model abstracts from the HRA component of New Option 1.

^b To compute the average share of spending out of pocket, we use the 2003 claims from all individuals in the baseline sample and apply each option's coverage details to this (common) sample. We then compute, for each option, the ratio of the resultant out-of-pocket expenses to the total claim amounts, and report the average (across employees in 2003) for each option. As a result, our computed average share of spending out of pocket abstracts from any differential behavioral effect of each contract. It does, however, account for the unmodeled small differences between the new and original options described above and in the text.

^c Premiums are normalized so that the lowest coverage is free for all employees. This is true in both the original and new options, up to small variation of several hundred dollars across employees. We report the average premium for employees in the baseline sample, pooling 2003 and 2004. Premiums vary by coverage tier; there is also some additional variation (across employees within coverage tier) in the incremental premiums associated with greater coverage options. The variation is based on the business unit to which each employee belongs (see Einav, Finkelstein, and Cullen 2010).

^d Statistics are based on all employee in the baseline sample, pooling 2003 and 2004.

One way to summarize the differences in consumer cost sharing under the different plans is to use the plan rules to simulate the average share of medical spending that would be paid out of pocket (counterfactually for most individuals) under different plans; we construct this measure of each plan's comprehensiveness using the spending of *all* 2003 employees and their realized medical claims, so that it does not reflect selection or moral hazard effects. Less generous plans correspond to those with higher consumer cost sharing. The results are summarized in the third row of each panel of Table 2. Combining the information on average enrollment shares of

lowest out-of-pocket maximum was \$2,000 (\$1,000) for nonsingle (single) coverage, corresponding to total annual spending of \$20,000 (\$10,000). Using the same realized claims distribution, we calculate that about 5.5 percent of employees would hit this out-of-pocket maximum.

the different plans with our calculation of the average cost sharing in the different plans, we estimate that, holding spending behavior constant, the change from the original options to the new options would have, on average, more than doubled the share of spending paid out of pocket, from about 13 to 28 percent.¹⁵

The plan descriptions in Table 2, and the subsequent parameterization of our model in Section IV, abstract from some additional details. First, while we model all plans as having a 10 percent in-network consumer coinsurance after the plan deductible is reached for all care, under the old options doctor visits and emergency room visits in fact had copays rather than coinsurance.¹⁶ Second, we have summarized (and modeled) the in-network features only. All of the plans have higher (less generous) consumer cost sharing for care consumed out of network rather than in network. We choose to model only the in-network rules (where more than 95 percent of spending occurs) in order to avoid having to model the decision to go in or out of network. Third, while in general the new options were designed to have higher consumer cost sharing, a wider set of preventive care services (including regular physicals, screenings, and well baby care) were covered with no consumer cost sharing under the new options; these preventive services account for less than 2 percent of medical spending in our sample. Finally, the least comprehensive of the new options (option 1) includes a health reimbursement account (HRA) into which the employer makes tax-free contributions that the employee can draw on to pay for out-of-pocket medical expenses, or roll over for subsequent years. In online Appendix F we explore specifications that try to account for these distinctive features of this option.

Table 3 shows plan transitions for employees who were in the old options in both 2003 and 2004 and for employees who were switched from the old to the new options in 2004. Two main features emerge. First, almost all employees under the old options in both years maintain the same coverage, which is to be expected given that the options and their prices did not change (but could also be driven by inertia in plan choices). Second, for those who get switched to the new options in 2004, there is far from a perfect correlation in the rank ordering of their choices under the old and new options. Over 40 percent of individuals move from the highest possible coverage under the old option to something other than the highest possible coverage under the new options, or vice versa. This is consistent with individuals making “active” choices under the new options, as suggested earlier.

III. Descriptive Evidence of Moral Hazard

We start by presenting some basic descriptive evidence of moral hazard in our setting, where by “moral hazard” we refer to the incremental medical spending associated with greater coverage, as defined in Section I. The analysis provides a feel for the basic identification strategy for moral hazard, as well as some suggestive evidence of heterogeneity in moral hazard and selection on it. At the same time,

¹⁵These numbers are based on the average out of pocket shares by plan calculated in Table 2 and the plan shares for the 2003–2006 sample (not shown). Using the 2003–2004 sample’s plan shares (shown in Table 2), we estimate that the move to the new options would, on average, raise the average out-of-pocket share from 12 to 25 percent.

¹⁶Specifically, they had doctor and ER copays of \$15 and \$75, respectively, or \$10 and \$50, depending on the plan. In practice, given the average costs of a doctor visit (\$115) and an ER visit (\$730) in our data, the switch from the copay to coinsurance did not make much difference for predicted out-of-pocket spending.

TABLE 3—PLAN TRANSITIONS (*Percent*)

	Old options in 2004	
	Highest coverage	All other coverages
2003		
Highest coverage	40.0	0.5
All other coverages	0.6	58.9
	New options in 2004	
	Highest coverage	All other coverages
2003		
Highest coverage	32.0	15.8
All other coverages	27.8	24.5

Notes: The table shows transition matrices across plan options for those in the old options in both 2003 and 2004 (top panel) and those who are switched to the new options in 2004 (bottom panel). Under the original options, the highest coverage is option 3. Under the new options, the highest coverage is option 5. See Table 2 for coverage details. The sample is limited to the 6,186 employees (82 percent of the baseline sample) who are in the data in both 2003 and 2004.

TABLE 4—SPENDING PATTERNS BY COVERAGE LEVEL

	Single coverage			Nonsingle coverage		
	Count	Mean	Median	Count	Mean	Median
Original plan options						
Highest coverage	512	3,130	557	2,318	6,634	2,670
All other coverages	1,031	1,795	233	3,035	5,768	2,288
New plan options						
Highest coverage	62	1,650	447	375	6,858	2,630
All other coverages	73	560	52	164	3,405	1,481

Notes: The table shows (contemporaneous) spending by coverage choice. Under the original options, the highest coverage is option 3. Under the new options, the highest coverage is option 5. See Table 2 for coverage details.

our descriptive exercise points to the difficulty in identifying heterogeneity in moral hazard and selection on it without a formal model of plan selection. The suggestive evidence as well as its important limitations together motivate our subsequent modeling exercise, which we turn to in the next section.

Average Moral Hazard.—We start with the (easier) empirical task of documenting the existence of some form of asymmetric information in our data. Table 4 reports realized medical spending as a function of insurance coverage in our baseline sample. The analysis—which is in the spirit of the “positive correlation test” in Chiappori and Salanie (2000)—shows that under either the old or new options individuals who choose more comprehensive coverage have systematically higher (contemporaneous) spending. This is consistent with the presence of adverse selection and/or moral hazard in our data.

To identify moral hazard separately from adverse selection, we take advantage of the variation in the option set faced by different groups of employees. Table 5 presents this basic difference-in-differences evidence of moral hazard for our baseline

TABLE 5—BASIC DIFFERENCE-IN-DIFFERENCES IN BASELINE SAMPLE

	Observations	Mean	Fraction with zero spending	10th	25th	50th	75th	90th
Control (switched after 2004)								
2003 spending	3,313	5,300	0.09	52	426	1,775	5,178	11,984
2004 spending	2,901	5,250	0.09	55	517	1,889	5,589	12,253
Treated (switched in 2004)								
2003 spending	682	5,201	0.08	79	581	1,957	5,048	12,644
2004 spending	674	4,856	0.10	22	447	1,610	4,622	9,467
Treated-Control differences (levels)								
2003 spending		-99	-0.01	27	155	182	-130	660
2004 spending		-394	0.01	-33	-70	-279	-967	-2,786
2004–2003 Difference (levels)								
Control		-50	0.00	3	91	114	411	269
Treated		-345	0.02	-57	-134	-347	-426	-3,177
Difference-in-differences (levels)								
		-295	0.02	-60	-225	-461	-837	-3,446
Difference (percentages)								
Control		-0.9%	0.0%	5.8%	21.4%	6.4%	7.9%	2.2%
Treated		-6.6%	22.0%	-72.2%	-23.1%	-17.7%	-8.4%	-25.1%
Difference-in-differences (percentages)								
		-5.7%	22.0%	-77.9%	-44.4%	-24.2%	-16.4%	-27.4%

sample. Specifically, we show various moments of the spending distribution in 2003 and in 2004 for the control group (employees who are covered by the old options in both years) and the treatment group (employees who are switched to the new options in 2004). The results show a strikingly consistent pattern across all the various moments of the spending distribution: spending falls for the treatment group, and tends to increase slightly for the control group.

Table 6 summarizes our central difference-in-differences estimates.¹⁷ Columns 1–3 show the results for our baseline 2003–2004 sample. The first column shows the difference-in-differences estimate when the dependent variable is measured in dollars, while columns 2 and 3 investigate specifications that give rise to a proportional moral hazard effect. Given the large fraction of employees with zero spending, we cannot estimate the model in simple logs. Instead, in column 2 we report estimates from a specification in which spending, m , is measured by $\log(1 + m)$,¹⁸ and column 3 reports a quasi-maximum likelihood Poisson model.¹⁹ The results suggest that the move to the new options is associated with an economically significant decline in spending.

An important concern about the results in columns 1–3 is that they are not very precise. This is reflected in the large standard errors of the estimates, and in the relatively large differences in the quantitative implications of the different point estimates. This lack of precision is driven by the fact that only about one-fifth of the employees in our

¹⁷ In online Appendix B we explore the sensitivity of our estimates to controlling for observable differences across employees, and investigate the validity of the identifying assumption underlying the difference-in-differences strategy.

¹⁸ Given that almost all individuals spend at least several hundred dollars (Figure 1), the results are not sensitive to the choice of 1 relative to some other small numbers. For the same reason, the estimated coefficients can be approximately interpreted as elasticities.

¹⁹ The quasi-maximum likelihood (QMLE)-Poisson model requires only that the conditional mean be correctly specified for the estimates to be consistent. See, e.g., Wooldridge (2002) for more discussion.

TABLE 6—DIFFERENCE-IN-DIFFERENCES ESTIMATES OF IMPACT OF CHANGE IN HEALTH INSURANCE OPTIONS ON ANNUAL MEDICAL SPENDING

	2003–2004 sample			2003–2006 sample		
	OLS in levels (1)	OLS in logs (2)	QMLE- Poisson (3)	OLS in levels (4)	OLS in logs (5)	QMLE- Poisson (6)
Estimated treatment effect	–297.2 (753.7) [0.70]	–0.35 (0.19) [0.08]	–0.06 (0.15) [0.69]	–591.8 (264.2) [0.034]	–0.175 (0.12) [0.17]	–0.114 (0.048) [0.018]
Mean dependent variable	5,232	6.91	5,232	5,392	6.9	5,392
Observations	7,570	7,570	7,570	14,638	14,638	14,638

Notes: The table shows the difference-in-differences estimate of the spending reduction associated with moving from the old options to the new options. The unit of observation is an employee-year. Dependent variable is the total annual medical spending for each employee and any covered dependents (or log of 1 + total spending in column 2 and column 5). The coefficient shown is the coefficient on an indicator variable that is equal to 1 if the employee's treatment group is offered the new health insurance options that year, and 0 otherwise. All regressions include year and treatment group fixed effects. We classify employees into one of four possible treatment groups—switched in 2004, switched in 2005, switched in 2006, or switched later—based on his union affiliation, which determines the year in which he is switched to the new health insurance options. Estimation is either by OLS or QMLE Poisson as indicated in the column headings. Standard errors (in parentheses) are adjusted for an arbitrary variance-covariance matrix within each of the 28 unions; *p*-values are in square brackets. Columns 1–3 show estimates for the 2003–2004 sample; columns 4–6 expand the sample to include 2003–2006.

sample are switched to the new benefits in 2004 (Table 1, top row). Therefore, in columns 4–6 we report analogous estimates from the 2003–2006 sample, during which more than half of the employees switched to the new benefits. As expected, the standard error of our estimates decreases substantially, and the quantitative implications of the results become much more stable across specifications. The estimated spending reduction is now statistically significant at the 5 percent level, with the point estimates suggesting a reduction of spending of about \$600 (column 4) or 11–17 percent (columns 5 and 6). In online Appendix B we show that the reduction in spending appears to arise entirely through reduced doctor and outpatient spending, with no evidence of a discernible effect on inpatient spending.

We can compute a back-of-the-envelope elasticity of health spending with respect to the out-of-pocket cost sharing by combining these estimates of the spending reduction with the average cost sharing of different plans (holding behavior constant). Given the distribution of employees across the different plans, the numbers in Table 2 suggest that the change from the old options to the new options should increase the average share of out-of-pocket spending from 12.6 to 28.4 percent in the 2003–2006 sample. Combining the point estimate of a \$591 reduction in spending (Table 6, column 4) with our calculation of the increase in cost sharing, our estimates imply an arc elasticity of medical spending with respect to out-of-pocket cost sharing of about -0.14 .²⁰ This is broadly similar to the widely used RAND experiment arc elasticity of medical spending of -0.2 (Manning et al. 1987; Keeler and Rolph 1988). Subsequent studies that have used quasi-experimental variation in

²⁰We compute an arc elasticity in which the proportional change in spending (and in consumer cost sharing) is calculated relative to the average observed across the old and new options, so that our results are more directly comparable with the existing literature. The arc elasticity is calculated as $[(q_2 - q_1)/(q_1 + q_2)/2]/[(p_2 - p_1)/(p_1 + p_2)/2]$ where *p* denotes the average consumer cost sharing rate. For the 2003–2006 sample, the proportional change in spending and cost sharing is 11 percent and 77 percent, respectively.

health insurance plans have tended to estimate elasticities of medical spending in the range of -0.1 to -0.4 .²¹

Heterogeneity in and Selection on Moral Hazard.—A necessary (but not sufficient) condition for selection on moral hazard is that there is heterogeneity in individuals' responsiveness to consumer cost sharing. To our knowledge, the experimental and quasi-experimental literature in health economics analyzing the impact of higher consumer cost sharing on spending has focused on average effects and largely ignored potential heterogeneity. This may in part reflect the fact that because health realizations are, by their nature, partially random, testing for heterogeneity in moral hazard is not trivial. It is particularly challenging without an explicit model of the nature of moral hazard that can, for example, provide guidance as to whether the effect of consumer cost sharing is additive or multiplicative.²² In addition, the typical nonlinear nature of health insurance coverage leads to heterogeneity in the intensity of the treatment, making it difficult to identify heterogenous effects from heterogenous treatments. In our specific context, a further subtlety is that it is the menu of plan options that varies in a quasi-experimental fashion, rather than the plan itself, making the actual individual coverage endogenous. All of these considerations motivate our reliance of a specific model of moral hazard and plan choice, which provides the basis for the primary empirical analysis.

Nonetheless, in online Appendix C we endeavor to present some suggestive evidence of what might plausibly be heterogeneity in moral hazard in the data. For example, we report the difference-in-differences estimates separately for observably different groups of workers. While many of the estimates are quite imprecise, the results are suggestive of larger moral hazard effects for older workers relative to younger workers and for sicker workers relative to healthier workers, and perhaps also for female and lower-income workers relative to male and higher-income workers, respectively. While suggestive, this type of exercise also points to the limitations of inferring heterogeneity in moral hazard across individuals from such simple descriptive evidence. For example, because the change is in menus rather than in specific plans, the extent of the treatment is driven by the endogenous plan choice from within the menu of options.

In that online Appendix we also look for suggestive evidence of selection on moral hazard. The pure comparative static of the model we present in Section I is that individuals with a greater behavioral response to coverage will choose greater coverage. Some suggestive evidence of such patterns comes from comparing the estimated behavioral response between those who chose more versus less coverage under the original options. Consistent with selection on moral hazard, we estimate

²¹ See Chandra, Gruber, and McKnight (2010), who provide a recent review of some of this literature as well as one of the estimated elasticities.

²² Without such a model, a nonparametric test for whether there is heterogeneity in moral hazard effects is possible to construct when there is no choice in health insurance and an exogenous change in health insurance coverage. In this case, a nonparametric test can be developed by relying on the panel nature of the data and comparing the joint distribution (before and after the introduction of a new benefit) of the quantiles of medical spending for the treatment group relative to the control group; the change in individual's spending rank (i.e., the joint distribution of the quantiles of spending) in the control group provides an estimate of the variation in ranking across individuals in their spending to expect simply from the random nature of health realizations. When an endogenous plan choice is present (as in our setting), however, a nonparametric test for heterogeneity in moral hazard is more challenging.

a reduction in spending associated with the move from the old options to the new options that is more than twice as large for those who originally had more coverage than for those who originally had less coverage, even though the reduction in cost sharing associated with the change in options (i.e., the treatment) is substantially larger for those who had less coverage. Yet the estimates are not precise, and, absent a model, it is difficult to separate the behavioral response from the endogenous plan choice from among the new options.

IV. Econometric Specification

A. Parameterization

We now turn to specify a more complete econometric model that is based on the economic model of individual coverage choice and utilization developed in Section I. This will allow us to jointly estimate coverage choices and utilization, relate the estimated parameters of the model to underlying economic objects of interest, and quantify how spending and welfare may be affected under various counterfactuals. The additional modeling assumptions in this section are of two different natures. First, we will need to specify more parametrically some of the objects introduced earlier (e.g., individuals' beliefs $F_\lambda(\cdot)$). Second, we need to specify what form of heterogeneity we allow across individuals, and for a given individual over time.

Our unit of observation is an employee i , in a given year t . We abstract from the specifics of the timing and nature of claims, and, as we have done so far, simply code utilization m_{it} as the total medical spending (in dollars) for the entire year. The individual faces the choice set of either the original plan options or the new plan options (as described in Table 2), depending on the year and the employee's union affiliation, which dictates whether and when he was switched to the new benefits options.

Using the model of Section I, recall that individuals are defined by three objects: their beliefs about their subsequent health status $F_\lambda(\cdot)$, their moral hazard parameter ω , and their risk aversion ψ . We assume that ω_i and ψ_i may vary across employees, but are constant for a given employee over time. It is the potential heterogeneity in ω_i that is the focus of the paper. We also assume that $F_\lambda(\cdot)$ is a (shifted) lognormal distribution with parameters $\mu_{\lambda, it}$, $\sigma_{\lambda, it}$, with support $(\kappa_{\lambda, it}, \infty)$, as explained below. That is, beliefs about health also vary across employees, and we allow $\mu_{\lambda, it}$ to be time varying to reflect the possibility that information about one's health evolves with time.

At the time of coverage choice individuals believe that

$$(6) \quad \log(\lambda_{it} - \kappa_{\lambda, it}) \sim N(\mu_{\lambda, it}, \sigma_{\lambda, it}^2),$$

and these beliefs are correct. Assuming a lognormal distribution for λ is natural, as the distribution of annual health expenditures is highly skewed (Figure 1). The additional parameter $\kappa_{\lambda, it}$ is used in order to capture the significant fraction of individuals who have no spending over an entire year. When $\kappa_{\lambda, it}$ is negative, the support of the implied distribution of λ_{it} is expanded, allowing for λ_{it} to obtain negative values, which in turn implies (when ω_i is not too large) zero spending. The parameter $\sigma_{\lambda, it}$ indicates the precision of the individual's information about his subsequent health.

It is the heterogeneity in $\mu_{\lambda, it}$, $\sigma_{\lambda, i}$, and $\kappa_{\lambda, i}$ that gives rise to the traditional form of adverse selection on the basis of expected health; i.e., on the basis of expected λ (denoted $\bar{\lambda}$), which is given by

$$(7) \quad \bar{\lambda}(\mu_{\lambda}, \sigma_{\lambda}, \kappa_{\lambda}) = \exp\left(\mu_{\lambda} + \frac{1}{2} \sigma_{\lambda}^2\right) + \kappa_{\lambda}.$$

That is, higher $\mu_{\lambda, it}$, $\sigma_{\lambda, i}$, or $\kappa_{\lambda, i}$ are all associated with higher expected λ , which all else equal leads to greater expected medical spending and greater cost by the insurance provider. All else equal, individuals with higher $\mu_{\lambda, it}$, $\sigma_{\lambda, i}$, or $\kappa_{\lambda, i}$ also prefer to choose greater coverage, thus giving rise to adverse selection.

Let \mathbf{x}_{it} denote a vector of observables that are taken as given, and let $\bar{\mathbf{x}}_i$ denote their within-individual average. In order to link the latent variables to observables, we make several parametric assumptions. First, we assume that $\log \omega_i$, $\log \psi_i$, and $\bar{\mu}_{\lambda, i}$ (which denotes the average (over time) of $\mu_{\lambda, it}$ for individual i) are drawn from a jointly normal distribution, such that²³

$$(8) \quad \begin{pmatrix} \bar{\mu}_{\lambda, i} \\ \log \omega_i \\ \log \psi_i \end{pmatrix} \sim N \left(\begin{pmatrix} \bar{\mathbf{x}}_i \boldsymbol{\beta}_{\lambda} \\ \bar{\mathbf{x}}_i \boldsymbol{\beta}_{\omega} \\ \bar{\mathbf{x}}_i \boldsymbol{\beta}_{\psi} \end{pmatrix}, \begin{pmatrix} \sigma_{\mu}^2 & \sigma_{\mu, \omega} & \sigma_{\mu, \psi} \\ \sigma_{\mu, \omega} & \sigma_{\omega}^2 & \sigma_{\omega, \psi} \\ \sigma_{\mu, \psi} & \sigma_{\omega, \psi} & \sigma_{\psi}^2 \end{pmatrix} \right).$$

We then assume a random effects structure on μ_{it} : we let μ_{it} vary over time, but assume that it is correlated within an employee, so that

$$(9) \quad \mu_{\lambda, it} = \bar{\mu}_{\lambda, i} + (\mathbf{x}_{it} - \bar{\mathbf{x}}_i) \boldsymbol{\beta}_{\lambda} + \epsilon_{\lambda, it},$$

where $\epsilon_{\lambda, it}$ is an i.i.d. normally distributed error term, with variance σ_{ϵ}^2 . The variance of $\mu_{\lambda, it}$ is then $\sigma_{\mu}^2 = \sigma_{\bar{\mu}}^2 + \sigma_{\epsilon}^2$. Finally, we assume that

$$(10) \quad \sigma_{\lambda, i}^{-2} \sim \Gamma(\gamma_1, \gamma_2) \mathbf{1} \{ \sigma_{\lambda, i}^2 \leq \bar{\sigma}^2 \},$$

and that

$$(11) \quad \kappa_{\lambda, i} \sim N(\bar{\mathbf{x}}_i \boldsymbol{\beta}_{\kappa}, \sigma_{\kappa}^2).$$

That is, $\sigma_{\lambda, i}^2$ is drawn from a right-truncated inverse gamma distribution, and $\kappa_{\lambda, i}$ is drawn from a normal distribution, and both are drawn independently from the other latent variables.

Thus, overall we estimate four vectors of mean shifters ($\boldsymbol{\beta}_{\lambda}, \boldsymbol{\beta}_{\omega}, \boldsymbol{\beta}_{\psi}, \boldsymbol{\beta}_{\kappa}$), eight variance and covariance parameters ($\sigma_{\mu}, \sigma_{\epsilon}, \sigma_{\omega}, \sigma_{\psi}, \sigma_{\kappa}, \sigma_{\mu, \omega}, \sigma_{\mu, \psi}, \sigma_{\omega, \psi}$), and two additional parameters (γ_1, γ_2) that determine the distribution of $\sigma_{\lambda, i}^{-2}$. Of course, an important decision is what observables $\bar{\mathbf{x}}_i$ shift which primitive, and whether we would like any observables to be excluded from one or more of the (four) equations.

²³For notational simplicity we consider $\bar{\mathbf{x}}_i$ to be the super-set of covariates, and implicitly assume some coefficient restrictions if we allow for different mean shifters for different latent variables.

To pay particular attention to the underlying variation emphasized in Section II, in all the specifications we experiment with, we include in \bar{x}_i treatment group fixed effects for each of the four treatment groups (see Table 1), as well as a year fixed effect on $\mu_{\lambda, it}$, the only time-varying latent variable. We also include coverage tier fixed effects since both the choice sets and spending varies substantially by coverage tier (see Table 2 and Figure 1, respectively), and a rich set of demographics, specifically age, gender, job tenure, income, and health risk scores.

B. Estimation

We estimate the model using Markov Chain Monte Carlo Gibbs sampling. The multidimensional unobserved heterogeneity naturally lends itself to such methods, as the iterative sampling allows us to avoid evaluating multidimensional integrals numerically, which is computationally cumbersome. The key observation is that the model we developed is sufficiently flexible so that we can augment the latent variables into the model and formulate a hierarchical statistical model. To see this, let $\theta_1 = \{\beta_\lambda, \beta_\omega, \beta_\psi, \beta_\kappa; \sigma_\mu, \sigma_\varepsilon, \sigma_\omega, \sigma_\psi, \sigma_\kappa, \sigma_{\mu, \omega}, \sigma_{\mu, \psi}, \sigma_{\omega, \psi}; \gamma_1, \gamma_2\}$ be the set of parameters we are interested in, and let $\theta_2 = \{\lambda_{it}, \mu_{\lambda, it}, \sigma_{\lambda, i}, \kappa_{\lambda, i}, \omega_i, \psi_i\}_{i=1, t=2003}^{i=N, t=2004}$ be the set of employee-year latent variables. The model is set up so that, even conditional on θ_1 , we can always rationalize the observed data—namely, plan choice and medical utilization—by appropriately finding a set of latent variables for each individual, θ_2 .

Thus, the iterative procedure is straightforward. We can first sample from the distribution of θ_1 conditional on θ_2 . Because, conditional on θ_2 , there is no additional information in the data about θ_1 , this part of the sampling is simple and quite standard. Then, we can sample from the distribution of θ_2 conditional on θ_1 and the information available in the data. This latter step is of course more customized toward our specific model, but does not introduce any conceptual difficulties. The full sampling procedure, the specific prior distributions we impose, and the resultant posteriors are described in detail in online Appendix D. We verified using Monte Carlo simulations that the procedure works effectively, and is robust to initial values. For our baseline results, the estimation appears to converge after about 5,000 iterations of the Gibbs sampler, so we drop the first 10,000 draws and use the last 10,000 draws of each parameter to report our results. The results we report are based on the posterior mean and posterior standard deviation from these 10,000 draws.

One important difficulty that our model introduces is related to our decision to not allow for an additive separable plan-specific error term. It is extremely common in applications of discrete choice (such as ours) to add such error terms, and often to assume that they are distributed i.i.d. across plans and individuals. Such error terms serve two important roles. First, they allow the researcher to rationalize any choice observed in the data through a large enough error term. Second, their independence makes the objective function of any M-estimator smooth, which is computationally attractive for numerical optimization. In the context of our application, however, we view such error terms as economically unappealing. The options from which individuals in our sample choose are financially rankable and are identical in their nonfinancial features. This makes one wonder what such error terms would capture that is outside of our model. The clear ranking of the options also makes the i.i.d.

nature of the error terms not very appealing. Instead, we introduce a fair amount of heterogeneity along the other dimensions of our model. Some of this heterogeneity (e.g., the heterogeneity in $\sigma_{\lambda,i}$ and $\kappa_{\lambda,i}$) is richer than the minimum required to capture the key economic forces we would like to capture, but this richness is what allows us to rationalize all observed choices in the data. This still leads to a model that is not very attractive for numerical optimization, which is one important reason why we use Gibbs sampling.²⁴

C. Identification

We now discuss the identification of the model. Conditional on the individual-behavior model described in Section I, the object of interest that we seek to identify is the joint distribution of $F_{\lambda}(\cdot)$, ω , and ψ . We have data on individuals' health insurance options, choices, and medical spending. Throughout the paper we make the strong assumption that individuals' beliefs about their subsequent health status ($F_{\lambda}(\cdot)$) are correct.²⁵

The model and its identification share many properties with some of our earlier work on insurance (Cohen and Einav 2007; Einav, Finkelstein, and Schrimpf 2010). The key novel element is that we now allow for moral hazard, and heterogeneity in it. The panel structure of the data and the staggered timing of the introduction of the new options are key in allowing us to identify this new element. We start our discussion of identification by considering nonparametric identification of our model with ideal data. We then discuss the ways in which our actual data is different from the ideal, thus requiring us to make additional parametric assumptions that aid in identification.

The two features of our dataset that are instrumental for identification are the panel structure of the data and the exogenous change in the health insurance options available to employees. In the ideal setting, we consider a case in which we observe individuals for a sufficiently long period before and a sufficiently long period after the change in coverage. Moreover, we assume that the choice set from which employees can choose coverage is continuous (for example, one can imagine a continuous coinsurance rate, and an increasing and differentiable mapping from coinsurance rate to premium).

In such a setting, our model is nonparametrically identified. To see this, note that such data provide us with two medical expenditure distributions, $G_i^{before}(m)$ and $G_i^{after}(m)$, for each individual i . Using the realized utility model (during the second period of the model), these two distributions allow us to recover for each individual $F_{i,\lambda}(\cdot)$ and ω_i . To see this, recall that abstracting from the truncation of medical spending at zero, our model implies that medical expenditure m_{it} is equal to $\lambda_{it} + \omega_i(1 - c_t)$.

²⁴In addition to our previous work (Cohen and Einav 2007; Einav, Finkelstein, and Schrimpf 2010), several other papers have estimated a discrete choice model without an i.i.d. error, for similar reasons. These include Keane and Moffitt (1998), Berry and Pakes (2007), and Goettler and Clay (2011).

²⁵While it is reasonable to question this assumption, absent direct data on beliefs, some assumption about beliefs is essential for identification. Otherwise, it is not possible to distinguish beliefs from other preferences that only affect choices, such as risk aversion. (See Einav, Finkelstein, and Schrimpf 2010 for a more detailed discussion of this point.) While we could instead assume some other (prespecified) form of biased beliefs, correct beliefs seem like a natural starting point.

If $F_{i,\lambda}(\cdot)$ is stable over time,²⁶ one can regress (for each employee i separately) m_{it} on a dummy variable that is equal to 1 after the change. The estimated coefficient on the dummy variable would then be an estimate of $\omega_i(c_{after} - c_{before})$, providing an estimate of ω_i . The distribution of λ_{it} can then be recovered by observing that $\lambda_{it} = m_{it} - \omega_i(1 - c_i)$, which is known.

Conditional on $F_{i,\lambda}(\cdot)$ and ω_i , individual i 's choice from a continuous set of options provides a unique mapping from choices to his coefficient of absolute risk aversion since—conditional on $F_{i,\lambda}(\cdot)$ and ω_i —the coefficient of risk aversion is the only unknown primitive that may shift employees' choices, and it does so monotonically. Thus, using information about $F_{i,\lambda}(\cdot)$ and ω_i and individual i 's choice from the continuous option set,²⁷ we can recover ψ_i . Since we recovered $F_{i,\lambda}(\cdot)$, ω_i , and ψ_i for each employee, we can now combine these estimates for our entire sample, and obtain the joint distributions of $F_{\lambda}(\cdot)$, ω , and ψ .

Our actual data depart from the ideal data described above in two main ways. First, although we have a panel structure, we only observe individuals for two periods in the baseline sample (that is limited to 2003 and 2004). Second, the choice set is highly discrete (including three to five options) rather than continuous. We thus make additional parametric assumptions to aid us in identification. This implies that our identification in the actual estimation cannot rely anymore on identifying the individual-specific parameters employee by employee. Rather, we observe a distribution of medical expenditures before the change and a distribution for medical expenditure after the change. We then identify the model by comparing the distribution after to the distribution before, using the untreated individuals to account for time-varying effects in medical spending, just like in the difference-in-differences analysis of Section III. Once the distribution of moral hazard, ω_i , is known, the remaining identification challenge is very similar to our earlier work mentioned above. In the working paper version (Einav et al. 2011) we provide a more detailed intuition for these last steps.

V. Results

A. Parameter Estimates

Table 7A presents the estimated parameters from estimating the model on the baseline sample of 7,570 employee-years. The top panel presents the estimated coefficients on the mean shifters of the four latent variables: $\mu_{\lambda, it}$ and $\kappa_{\lambda, i}$, which affect expected health risk ($E(\lambda_{it})$), ω_i , which affects moral hazard, and ψ_i , which captures risk aversion. The middle panel report the estimated variance-covariance matrix and the bottom panel reports the estimates of the rest of the parameters. In Table 7B we report some implied quantities of interest that are derived from the estimates. The latter may be more easy to interpret, so we focus much of the discussion on them.

²⁶ If $F_{i,\lambda}(\cdot)$ changes over time, one could parameterize, identify, and estimate the autocorrelation structure with a sufficiently long panel (the health risk score variable, which varies over time for a given individual, is quite useful in this regard). We therefore treat $F_{i,\lambda}(\cdot)$ as stable over time throughout this section.

²⁷ This can be done using either the options set before the change or after. In fact, the ideal data leads to over-identification, so it could allow us to test for the model's assumptions and/or to enrich the model.

TABLE 7A—PARAMETER ESTIMATES

	μ_λ (Health risk)	κ_λ (Health risk)	$\ln(\omega)$ (Moral hazard)	$\ln(\psi)$ (Risk aversion)
<i>Mean shifters</i>				
Constant	6.11 (0.14)	-389 (73)	5.31 (0.24)	-5.57 (0.10)
Coverage tier				
Single	(Omitted)	(Omitted)	(Omitted)	(Omitted)
Family	0.19 (0.08)	57 (51)	-0.58 (0.18)	-0.88 (0.07)
Emp+Spouse	0.27 (0.09)	44 (53)	-0.66 (0.22)	-0.95 (0.07)
Emp+Children	0.24 (0.08)	185 (47)	-0.28 (0.21)	-0.91 (0.06)
Treatment group				
Switch 2004	-0.01 (0.07)	-278 (43)	-0.24 (0.11)	-0.31 (0.05)
Switch 2005	-0.10 (0.06)	-78 (38)	0.07 (0.12)	-0.23 (0.05)
Switch 2006	0.12 (0.07)	-94 (37)	0.01 (0.12)	-0.07 (0.05)
Switch later	(Omitted)	(Omitted)	(Omitted)	(Omitted)
Demographics				
Age	-0.01 (0.003)	-5 (1.8)	-0.01 (0.006)	0.01 (0.002)
Female	0.18 (0.08)	94 (39)	-0.08 (0.13)	-0.07 (0.06)
Job Tenure	0.002 (0.003)	-2.3 (1.6)	0.002 (0.004)	0.003 (0.002)
Income	0.003 (0.002)	6 (0.9)	0.001 (0.003)	-0.0003 (0.001)
Health risk score				
1st quartile (< 1.119)	(Omitted)	(Omitted)	(Omitted)	(Omitted)
2nd quartile (1.119 to 1.863)	0.91 (0.07)	305 (59)	0.13 (0.29)	-0.41 (0.06)
3rd quartile (1.863 to 2.834)	1.48 (0.08)	242 (81)	1.79 (0.27)	-0.66 (0.06)
4th quartile (> 2.834)	2.05 (0.09)	-416 (120)	3.38 (0.22)	-0.89 (0.07)
2004 Time dummy	-0.12 (0.02)	—	—	—
<i>Variance-covariance matrix</i>				
$\bar{\mu}_\lambda$	0.20 (0.03)	$\ln(\omega)$	$\ln(\psi)$	
$\ln(\omega)$	—	-0.03 (0.04)	-0.12 (0.02)	
$\ln(\psi)$	—	0.98 (0.08)	-0.01 (0.03)	
		—	0.25 (0.02)	
<i>Additional parameters</i>				
σ_μ	0.33 (0.03)			
σ_κ	290 (12)			
γ_1	0.04 (0.004)			
γ_2	15 (1.2)			

Notes: The table presents our baseline parameter estimates based on our baseline sample of 7,570 employees. As described in the text, the estimates are based on a Gibbs sampler; the table reports the posterior mean and the posterior standard deviations in parentheses. All time-varying demographics are set to their mean over the two years, except for health risk score, which we allow to be a time-varying shifter of μ_λ . “Treatment group” dummies refer to indicator variables based on the year your union’s benefits were switched to the new benefits. Higher risk scores correspond to worse predicted health.

TABLE 7B—IMPLIED QUANTITIES

	$E(\lambda)$	ω	ψ
<i>Unconditional statistics</i>			
Average	4,340 (200)	1,330 (59)	0.0019 (0.00002)
SD	5,130 (343)	3,190 (320)	0.0020 (0.00007)
<i>Unconditional correlations</i>			
$E(\lambda)$	1.00	0.24 (0.03)	-0.36 (0.01)
ω	—	1.00	-0.15 (0.01)
ψ	—	—	1.00
<i>Marginal effects on $E(\lambda)$</i>			
Coverage tier			
Single	(Omitted)		
Family	360 (75)		
Emp+Spouse	700 (85)		
Emp+Children	300 (69)		
Treatment group			
Switch 2004	260 (71)		
Switch 2005	-320 (67)		
Switch 2006	640 (70)		
Switch later	(Omitted)		
Demographics			
Age	-42 (3.6)		
Female	720 (71)		
Job tenure	5.9 (3.1)		
Income	13 (2)		
Health risk score			
1st quartile (< 1.119)	(Omitted)		
2nd quartile (1.119 to 1.863)	1,600 (96)		
3rd quartile (1.863 to 2.834)	4,000 (200)		
4th quartile (> 2.834)	8,500 (370)		
2004 Time dummy	-590 (29)		

Notes: The table reports some implied quantities of interest that are derived from the estimated parameters in Table 7A. Posterior standard deviations are in parentheses.

Overall, as shown in the top panel of Table 7B, the estimates imply an average health risk ($E(\lambda)$) of about \$4,340 per employee-year. We estimate an average moral hazard parameter (ω) that is about 30 percent of the average health risk, or about \$1,330; by way of context, recall that ω is approximately the size of the spending effect as we move individuals from no insurance to full insurance (see equation (3)).²⁸

We estimate statistically significant and economically large heterogeneity in each one of the components: health, moral hazard, and risk aversion. One way to gauge the magnitude of this heterogeneity is in the top panel of Table 7B. Our estimates indicate a standard deviation for expected health risk ($E(\lambda)$) of about \$5,100, or a coefficient of variation of about 1.2; the standard deviation of realized health (λ) is, not surprisingly, much larger at \$25,000 (not shown). Moral hazard (ω) is also estimated to be highly heterogenous, with a standard deviation across employees of about \$3,200, or a coefficient of variation that is greater than 2. Finally, we estimate a coefficient of variation for absolute risk aversion (ψ) that is about one.

The unconditional correlations (Table 7B, middle panel) are all statistically significant, and their signs seem reasonable. We estimate that the unconditional correlation between moral hazard (ω) and expected health risk ($E(\lambda)$) is positive and reasonably important (0.24). This likely reflects the fact that in our model moral hazard type (ω) is measured in absolute (dollar) terms rather than relative to health, so individuals with higher $E(\lambda)$ (i.e., worse health) have more opportunities to exercise moral hazard. The correlation between risk aversion and health risk (and moral hazard) is negative, perhaps reflecting the fact that individuals who are more risk averse are also those who take better care of their health. A similar pattern was documented by Finkelstein and McGarry (2006) in the context of long-term care insurance. Finally, as may be expected, we estimate a strong correlation in $\mu_{\lambda, it}$, over time, of 0.5 (not shown), suggesting that much of an individual's health risk is persistent over time, for example due to chronic conditions.

The signs of the covariates seem generally sensible. The bottom panel of Table 7B summarizes the effects of covariates on $E(\lambda)$ by combining their separate effects on μ_{λ} and κ_{λ} . As could be expected, the health risk scores are an important predictor of expected health risk $E(\lambda)$, shifting it by thousands of dollars in the expected direction. We also estimate that female employees and employees with nonsingle coverage are associated with worse expected health (higher $E(\lambda)$). One should interpret these latter effects cautiously, however, as health risk scores are partialled out and are highly correlated with these other variables. This may also explain why the residual effect of income and tenure on expected health appears negligible.

Our estimates also imply (Table 7A, top panel) that employees with higher (i.e., worse) health risk scores are associated with greater moral hazard and lower risk aversion. Again, this likely reflects our choice to model moral hazard in absolute

²⁸ We estimate an average coefficient of absolute risk aversion of about 0.0019, but caution against trying to compare this to existing estimates. In our model, realized utility is a function of both health risk and financial risk, while in other papers that estimate risk aversion from insurance choices (e.g., Cohen and Einav 2007, Handel 2012) realized utility is only over financial risk. Thus, the estimated level of risk aversion is not directly comparable; indeed, one could add a separable health-related component to utility that is affected only by λ to change the risk aversion estimates, without altering anything else in the model.

TABLE 8—MODEL FIT: CHOICE PROBABILITIES (*Percent*)

Plan	Data	Model
<i>Original options (N = 6,896)</i>		
Option 1	1.2	2.0
Option 2	58	57
Option 3	41	41
<i>New options (N = 674)</i>		
Option 1	5.9	5.0
Option 2	0.5	5.0
Option 3	1.9	1.0
Option 4	27	11
Option 5	65	76

Notes: The table reports the actual and predicted choice probabilities of each plan. Plans are numbered from lowest to highest coverage. For plan details see Table 2.

terms rather than relative to health. Conditional on health risk scores, employees with single coverage appear to be associated with greater moral hazard as well as with greater risk aversion. This may represent different process of decision making regarding health coverage and health care utilization when regarding oneself versus his family members.

B. Model Fit

In Table 8 we report the actual and predicted plan choice probabilities. We fit the choices of employees who are choosing from the original plan options remarkably well. The fit of the choices from the new options is also reasonable, but not as good as the fit for the original options. This is likely because there are many fewer employees in the baseline sample who are subject to the new options. Thus, to the extent that the same model attempts to rationalize the choices from both the old and new options, it is natural that more weight is given to trying to fit choices from the old menu, leading to slightly worse fit for those choosing from the new menu.

Figure 2 reports the actual and predicted distributions of medical expenditure. The top panel reports the fit for the individuals facing the old options, and the bottom panel reports the fit for the individuals facing the new options. Overall, the fit is quite reasonable. For example, the predicted average spending is within 10 percent of actual average spending under both the original and new options, and the medians also fit quite well. We tend to overpredict the fraction of individuals who have no spending under the new options, but this again is likely driven by the relatively small number of employees who are switched to the new options in our estimation sample.

Finally, we note that if we simulate data based on our parameter estimates and then run the difference-in-differences analysis we report on in Section IV, we predict about an 8 percent reduction in spending associated with moving from the old option set to the new option set. This is broadly similar to the difference-in-differences estimates we obtained for the actual data (Table 6, columns 1–3). Given how imprecise our difference-in-differences estimates are, however, both in the actual data and in the simulated data, we caution against making too much of any comparison. The lack of statistical significance of the difference-in-differences estimate in the

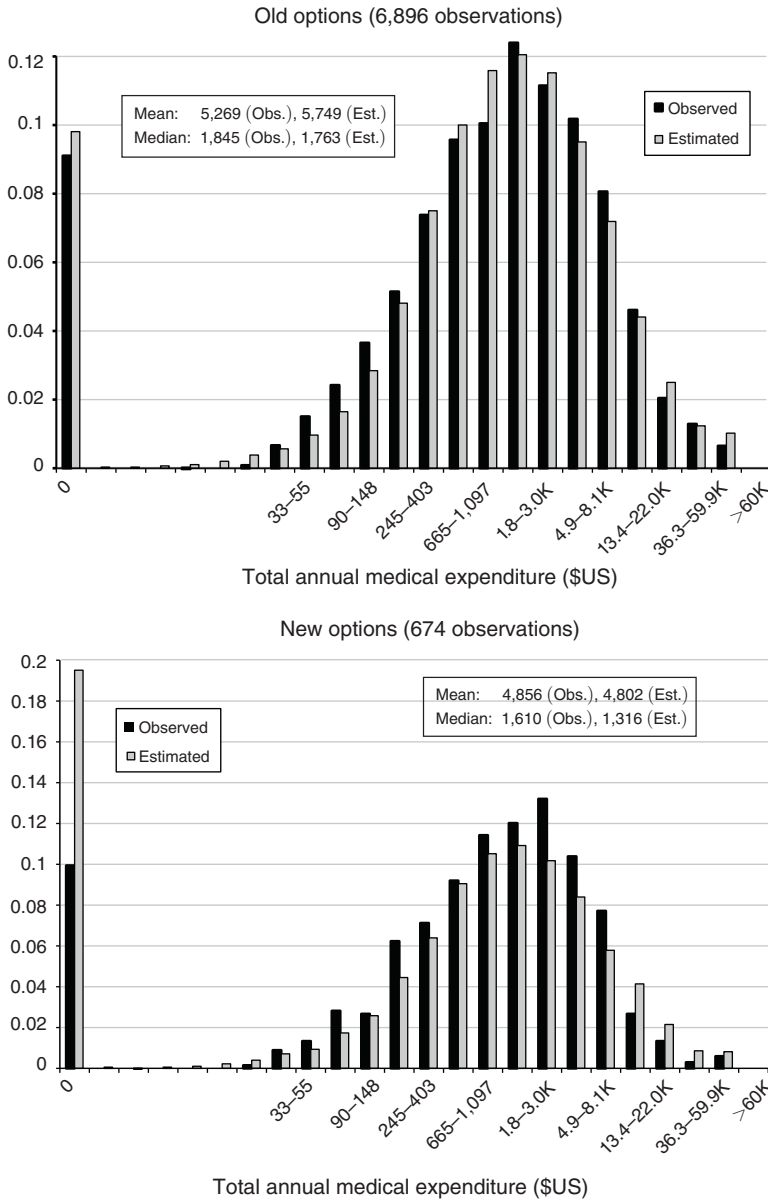


FIGURE 2. MODEL FIT: MEDICAL SPENDING DISTRIBUTIONS

Notes: The figure presents the distribution of total annual medical expenditure, in the data and in model simulations based on the estimated parameters. The graph uses a log scale, such that the second bin covers expenditure lower than $\exp(0.5)$, the next covers expenditures between $\exp(0.5)$ and $\exp(1)$, and so on; the x -axis labels show the corresponding dollar amounts of selected bins. The top panel compares spending of individuals who faced the original options, and the bottom panel compares the spending distribution of individuals who faced the new options.

simulated data, relative to the reasonably precise estimates of the model parameters, suggests that a more complete model of unobservable heterogeneity and endogenous plan choice is important in increasing precision.

TABLE 9—SPENDING IMPLICATIONS OF MORAL HAZARD ESTIMATES

	Mean	SD	10th	25th	50th	75th	90th
Spending difference as we move from no to high-deductible plan	348	749	0	0	48	316	1,028
Spending difference as we move from full to no insurance	1,273	3,181	0	86	310	1,126	3,236

Notes: The table reports the implied spending implications if we move different employees across plans. For each employee, we use the model estimates to compute his decline in expected annual expenditure as we change his insurance plan. In the top row, we move each employee from the highest coverage option under the new benefit options (option 5) to the lowest coverage option under the new benefit options (option 1); roughly speaking, this entails moving from a plan with no deductible to a plan with a high deductible; see Table 2 for more details. In the bottom row, we move each employee from full to no insurance. The table then summarizes the cross-sectional distribution of the spending effects. The estimates are primarily driven by the estimated distribution of ω , but they take into account the truncation of spending at zero by integrating over the conditional (on ω) distribution of λ .

C. Moral Hazard Estimates

The parameter ω_i captures moral hazard in our model. Recall that, abstracting from the truncation of spending at zero, employee i would spend λ_{it} in year t if he had no insurance, and with full insurance would spend $\lambda_{it} + \omega_i$. Thus, ω_i can be thought of as the scope for moral hazard. As discussed, the top panel of Table 7B reports that the estimated average of ω_i is about \$1,330, or about 30 percent of the estimated health risk (the average of λ_{it}).

Table 9 reports an alternative way one could quantify moral hazard. In the top row of the table, we calculate each employee's expected decline in medical expenditure as we move him from the highest to the lowest coverage in the new options. We will feature the move (or choice) between these two options in all of our subsequent counterfactual exercises. Recall that, as we have modeled these options, moving from the highest to the lowest coverage primarily entails moving someone from a plan with no deductible to a plan with a high deductible, specifically a \$3,000 deductible for nonsingle coverage, or \$1,500 for single coverage (Table 2). We estimate that the average spending effect from this move is \$348. The second row reports a similar exercise, but considers moving individuals from full insurance to no insurance. We estimate an average spending reduction of \$1,273; this is slightly lower than the average ω_i of \$1,330 reported earlier (Table 7B) precisely because of the truncation of spending at zero.

These economically meaningful estimates of moral hazard satisfy one necessary condition for selection on moral hazard—the focus of our paper—to be important. A second necessary condition is that moral hazard be heterogenous. Indeed, we find important heterogeneity in our moral hazard estimates across individuals. For example, the estimated variance of $\log(\omega)$ is about 1, and highly statistically significant (Table 7A), implying that an employee who is one standard deviation above the mean is associated with a moral hazard parameter that is almost three times greater than the mean, and an employee who is one standard deviation below the mean has a moral hazard parameter that is less than a half of the mean. As shown in the top of Table 7B, across individuals, the standard deviation of ω_i is almost \$3,200, and the coefficient of variation of ω is more than 2.

Again, Table 9 reports more economics-motivated measures of heterogeneity in moral hazard. The top row shows that the spending decline as we move individuals from the no-deductible plan to the high-deductible plan has a standard deviation of \$749, compared to the mean of \$348. The median spending reduction is only \$48, while the ninetieth percentile exhibits a spending reduction of more than \$1,000. Similarly, as we move individuals from full insurance to no insurance, we estimate that the median reduction in spending is \$310, but the 90th percentile of the spending reduction distribution is greater than \$3,000.²⁹

D. Selection on Moral Hazard

The fact that individuals are heterogenous in their moral hazard responses to coverage does not, of course, mean that they select on it in any quantitatively meaningful way. That is, it is conceivable that heterogeneity in other factors is more important in determining plan choice. As one way to gauge the quantitative importance of selection on moral hazard, we examine how the choice of coverage varies with the quantiles of the marginal distribution of moral hazard ω , and compare this to how the choice of coverage varies with the quantiles of the marginal distribution of risk aversion ψ , and of expected health risk $E(\lambda)$. Once again, we focus on the choice between the highest-coverage and lowest-coverage plans in the new options (see Table 2). Loosely, our exercise resembles the introduction of a high-deductible health insurance plan into a setting where previously there was only a no-deductible plan. We set the premiums so that, on average, 10 percent of our sample chooses the high-deductible plan.

Figure 3 reports the results. It shows the fraction of individuals choosing the high-deductible coverage, conditional on the individual being in each quantile of the marginal distribution of moral hazard ω , of risk aversion ψ , and of expected health risk $E(\lambda)$. We present two different sets of results. The top panel presents the pattern while taking as given the underlying correlation structure among these objects. This panel can be thought of as giving the empirical answer to the question of how much selection there is, on net, on each of the latent primitives that we model. Given the flexible correlations we allow for, these patterns are, a priori, of ambiguous sign. The bottom panel repeats the same exercise but “shuts down” the effect of the correlation structure. To do so, we compute the marginal distributions (unconditional on observables) of each of the three latent variables that affect plan choice (ω , ψ , and $E(\lambda)$), and draw values for the other two latent variables independent of the value of the variable for which the graph is drawn. This panel can be thought of as giving the answer to the conceptual comparative static exercise of how much selection there is on one latent factor, holding the other factors constant. As discussed previously, demand for higher coverage generally increases in expected health risk, in risk aversion, and in moral hazard. Our purpose here is to assess the relative magnitudes. Taken together, the two panels help inform not only whether, empirically, there is

²⁹ We explored whether our modeling of moral hazard as an additive shifter of spending is important in driving our estimates of significant heterogenous moral hazard. Online Appendix E shows results from an alternative model in which moral hazard is modeled as multiplicative rather than additive; we continue to find substantial heterogeneity in moral hazard in this alternative model.




FIGURE 3. SELECTION ON MORAL HAZARD RELATIVE TO OTHER SOURCES OF SELECTION

Notes: The figure illustrates the relative importance of the three different sources of selection that we model. We consider an individual's choice between two available options: the no-deductible and high-deductible plans among the new set of options (see Table 2, options 5 and 1, respectively). We assume the observed (averaged within each coverage tier) premiums for these two options. Each point in the figure indicates the fraction of individuals choosing the high-deductible (i.e., low-coverage) option relative to the no-deductible (high-coverage) option. We consider three sources of selection: $E(\lambda)$ (risk), ω (moral hazard), and ψ (risk aversion). For each of them, we compute the fraction choosing the high deductible at different quantiles of the distribution. In the top panel, we take into account the correlation between each component and the others, while in the bottom panel we repeat the same exercise but draw the other components of the model randomly from their marginal distribution (that is, assuming no correlation).

selection on moral hazard and of what sign (top panel) but also the extent to which any such selection is primarily “direct” selection based on moral hazard rather than “indirect” selection arising from the correlation structure between moral hazard and other factors that may be driving plan choice.

The results in the top panel indicate that empirically there is selection on moral hazard of the expected sign, with higher moral hazard types (higher ω) less likely to choose the high-deductible plan. In terms of the substantive importance of this selection, both panels reveal a similar qualitative pattern: selection on moral hazard is substantially larger than selection on risk aversion and of similar magnitude to selection on health risk. For example, the top panel indicates that moving from the tenth percentile to the ninetieth percentile of the moral hazard distribution is associated with about a 23 percentage point decline in the demand for the high-deductible plan, while moving from the tenth to the ninetieth percentile of the expected health risk distribution is associated with about a 24 percentage point decline in the demand for the high-deductible plan. While some of this reflects the underlying correlation structure, the “pure” comparative static shown in the bottom panel produces quite comparable magnitudes. This suggests that much of this selection on moral hazard is “direct” selection. In other words, in making plan choices, individuals select not only based on the expected levels of spending that they would incur with no insurance, but also on their expected slope, or incremental spending due to insurance.

By contrast, we find selection on risk aversion considerably less important than selection on either moral hazard or expected health. In our data (see Figure 3, panel A) there is very little variation in demand for the high-deductible plan across the centiles of the risk aversion distribution (reflecting various correlations), and even the “pure” comparative static (Figure 3, panel B) suggests only about a 15 percentage point range between the tenth and ninetieth percentiles.

E. Implications for Spending

We investigate the implications of the selection on moral hazard that we detect for attempts to combat moral hazard through higher consumer cost sharing. To this end, we perform counterfactual analyses of the spending reduction associated with introducing a lower coverage option. Given our finding that higher “moral hazard types” prefer greater coverage, accounting for this selection on moral hazard suggests that introducing plans with greater consumer cost sharing will produce less of a spending reduction than would be estimated if selection on moral hazard were ignored, and it were assumed that those who select the lower coverage option are drawn at random from the “moral hazard type” distribution.

In the health care sector, the impact of consumer cost sharing on moral hazard is an issue of considerable policy as well as academic interest. The size and rapid growth of the health care sector, and the pressure this places on public sector budgets, has created great interest among both academics and policymakers in possible approaches to reducing health care spending. Encouraging individuals to enroll in plans with higher consumer cost sharing, such as the tax-advantaged Health Savings Accounts designed to increase enrollment in high-deductible plans, is seen as one potentially promising approach to reducing health spending.

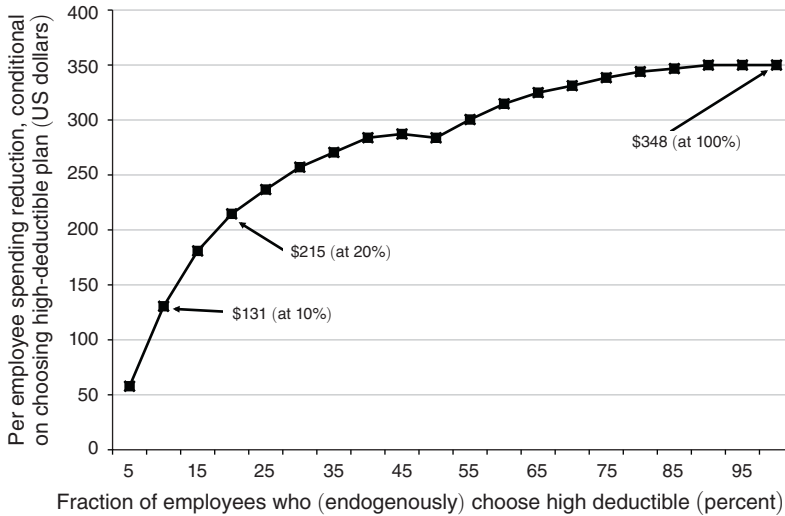


FIGURE 4. SPENDING IMPLICATIONS OF SELECTION ON MORAL HAZARD

Notes: The figure illustrates the potential spending implications arising from selection on moral hazard. To construct the figure, we use an exercise similar to the one used for Figure 3. For each individual, we use the model estimates to compute his decline in expected annual expenditure as we move him from the highest coverage (no deductible) to the lowest coverage (high deductible) in the new benefits options (see Table 2, options 5 and 1, respectively). We then vary the relative price of the highest coverage, allowing employees to choose endogenously between the two options, and report the per-employee expected decrease in spending for the group of individuals who choose the lowest coverage at each price. Without selection on moral hazard, the curve would have been flat. Selection on moral hazard implies that those with the lowest moral hazard effects of insurance are those who have the lowest willingness to pay for incremental coverage and are therefore the first (as the price of coverage increases) to switch from higher to lower coverage. *Ceteris paribus*, therefore, selection on moral hazard generates an upward-sloping curve; this can be offset through the correlation between moral hazard and other components of demand (such as risk aversion or health risk).

To examine the implications of selection on moral hazard for analysis of such efforts, Figure 4 engages in the same exercise as in Figure 3 of giving employees in our sample a choice between the no-deductible and high-deductible health insurance plans in the new options. In Figure 3 we fixed the price of each option and reported the fraction of each quantile of a latent variable who choose each plan. In Figure 4 we instead gradually increase the (relative) price of the higher-coverage (no-deductible) option, and assess how much selection is associated with the group of employees who endogenously select the lower coverage (high-deductible) option at each given price. To show the extent of selection, the figure reports the average per-employee decline in annual spending for those employees who endogenously select the high-deductible plan at each price.

The figure illustrates strong selection on moral hazard, especially when the share of the high-deductible plan is small. For example, when the price of the no-deductible coverage is low enough so that only 10 percent of the employees select the high-deductible coverage, the average (per-employee) spending decline for those who select the high-deductible plan instead of the no-deductible plan is just over \$130. By contrast, were all employees to choose the high-deductible plan instead of the

no-deductible plan, we estimate the per-employee spending decline would be about \$350. As noted in the introduction, the common practice in the literature on health insurance and moral hazard is to look for experimental variation that randomly moves individuals across plans. Such variation would recover the unconditional average effect of coverage (which is \$348 in our context); this does not account for selection on moral hazard and will therefore substantially overestimate the spending reduction associated with the introduction of the high-deductible plan when only a small share of individuals select it.

This selection reflects the earlier observation that, all else equal, individuals who are associated with higher moral hazard (higher ω_i) have higher willingness to pay for insurance, and are therefore the ones who would be the last to switch to the lowest coverage, as we gradually increase the price of highest coverage. It is somewhat interesting that in our setting the selection on moral hazard becomes less important (i.e., the slope of the line in Figure 4 becomes less steep) at higher levels of prices for the no-deductible plans (which leads to greater fractions choosing the high-deductible plan). The same underlying forces are still in play, but are offset by the correlation structure with other primitives.

VI. Illustrative Welfare Analysis

Our findings of selection on moral hazard also have implications for policies aimed at reducing selection. Analysis of how to mitigate selection often focuses on risk adjustment—whereby an individual’s insurance premiums are adjusted on the basis of individual covariates (such as age, gender, and prior health conditions) that are predictive of expected medical spending. From this perspective, the potential for selection on moral hazard suggests that investments in better monitoring technologies—such as coinsurance rates that vary across diagnoses (e.g., heart attack versus headache) or types of health care (e.g., doctor visits versus inpatient services) with different behavioral responsiveness to insurance—may also be effective at ameliorating adverse selection. Our final set of counterfactual analyses considers these issues of contract design by using our model to go further out-of-sample and analyze the impact of alternative contract designs on social welfare.³⁰

A. Measuring Welfare

Our standard measure of consumer welfare will be the notion of certainty equivalent. That is, for an individual defined by $(F_\lambda(\cdot), \omega, \psi)$, we denote the certainty

³⁰Our framework assumes that any moral hazard–induced expenditure represents pure waste from a societal perspective. While this seems (to us) a natural benchmark—and is in line with the traditional view of moral hazard as a distortion that arises from misaligned incentives that result from the fundamental (risk smoothing) properties of providing insurance (e.g., Holmström 1979)—it is reasonable to question whether this literal interpretation is appropriate as a normative statement in the context of health insurance. For example, in the absence of subsidies, liquidity-constrained, ill-informed, or myopic consumers may underconsume medical care. Absent any clear guidance as to the nature and magnitude of any such frictions, we abstract from them in our model and note that it is not necessary to interpret these welfare results in a normative light. It is still a useful metric by which we can quantify and compare the effects of moral hazard, adverse selection, and different contract design features.

equivalent to a contract j by the scalar e_j that solves $-\exp(-\psi e_j) = v_j(F_\lambda(\cdot), \omega, \psi)$, or

$$(12) \quad e_j(F_\lambda(\cdot), \omega, \psi) \equiv -\frac{1}{\psi} \ln \left[\int \exp(-\psi u^*(\lambda, \omega, j)) dF_\lambda(\lambda) \right].$$

Our assumption of CARA utility over (additively separable) income and health implies no income effects. Because y and p_j are taken as given (at the time of utilization choice), it will be convenient to define

$$(13) \quad \tilde{u}(m; \lambda, \omega, j) = \left[(m - \lambda) - \frac{1}{2\omega}(m - \lambda)^2 \right] - c_j(m),$$

so that (based on equation (1)) $u(m; \lambda, \omega, j) = \tilde{u}(m; \lambda, \omega, j) + y - p_j$. It will also be convenient to denote $\tilde{u}^*(\lambda, \omega, j) \equiv \tilde{u}(m^*(\lambda, \omega, j); \lambda, \omega, j)$.

To see the implications of no income effects, we can substitute $u^*(\lambda, \omega, j) = \tilde{u}^*(\lambda, \omega, j) + y - p_j$ into equation (12) and reorganize to obtain

$$(14) \quad e_j(F_\lambda(\cdot), \omega, \psi) \equiv \tilde{e}_j(F_\lambda(\cdot), \omega, \psi) + y - p_j \\ \equiv -\frac{1}{\psi} \ln \left[\int \exp(-\psi \tilde{u}^*(\lambda, \omega, j)) dF_\lambda(\lambda) \right] + y - p_j,$$

so that $\tilde{e}_j(F_\lambda(\cdot), \omega, \psi)$ captures the welfare from coverage, and residual income enters additively. Using this notation, differences in $\tilde{e}(\cdot)$ across contracts with different coverages capture the willingness to pay for coverage. For example, an individual defined by $(F_\lambda(\cdot), \omega, \psi)$ is willing to pay at most $\tilde{e}_k(F_\lambda(\cdot), \omega, \psi) - \tilde{e}_j(F_\lambda(\cdot), \omega, \psi)$ in order to increase his coverage from j to k .

We assume that insurance providers are risk neutral, so that the provider's welfare is given by his expected profits, or

$$(15) \quad \pi_j(F_\lambda(\cdot), \omega) \equiv p_j - \int [m^*(\lambda, \omega, j) - c_j(m^*(\lambda, \omega, j))] dF_\lambda(\lambda),$$

where the integrand captures the share of the utilization covered by the provider under contract j . Total surplus s_j is then given by

$$(16) \quad s_j(F_\lambda(\cdot), \omega, \psi) = e_j(F_\lambda(\cdot), \omega, \psi) + \pi_j(F_\lambda(\cdot), \omega) \\ = \tilde{e}_j(F_\lambda(\cdot), \omega, \psi) + y \\ - \int [m^*(\lambda, \omega, j) - c_j(m^*(\lambda, \omega, j))] dF_\lambda(\lambda).$$

That is, total surplus is simply certainty equivalent minus expected cost.

To gain intuition, it may be useful to discuss the nature of the efficient contract in this setting. Because of our CARA assumption, premiums are transfers that

do not affect total surplus. Therefore, the efficient contract can be characterized by the efficient coverage function $c^*(\cdot)$ that maximizes total surplus (as given by equation (16)) over the set of possible coverage functions. Such optimal contracts would trade off two offsetting forces. On one hand, an individual is risk averse while the provider is risk neutral, so optimal risk sharing implies full coverage, under which the individual is not exposed to risk. On the other hand, the presence of moral hazard makes an insured individual's privately optimal utilization choice socially inefficient; any positive insurance coverage makes the individual face a health care price that is lower than the social cost of health care, leading to excessive utilization. Efficient contracts will therefore resolve this trade-off by some form of partial coverage (Arrow 1971; Holmström 1979). For example, it is easy to see that no insurance ($c^*(m) = m$) is efficient if individuals are risk neutral or face no risk ($F_\lambda(\cdot)$ is degenerate), and that full insurance ($c^*(m) = 0$) is efficient when moral hazard is not present ($\omega = 0$). In all other situations, the efficient contract is some form of partial insurance.

B. Welfare Implications

Table 10 reports our welfare results. Once again we restrict our attention to a choice between the no-deductible and high-deductible plans under the new options (Table 2, options 5 and 1, respectively). Throughout this section we make the simplifying assumption of perfect competition for the incremental coverage among providers of the no-deductible plan, so that the incremental price of the no-deductible plan breaks even for those who provide it: incremental price is equal to incremental cost.³¹ We report the implications of various counterfactual contracts for the equilibrium (incremental) premium for the no-deductible plan, the share choosing this plan, expected spending per employee, and total welfare (or surplus) per employee. Our primary focus is on the consequence of different contract designs for total welfare (i.e., the sum of consumer welfare and producer welfare), which in our context is the certainty equivalent minus expected costs (see equation (16)).

The first row presents the “status quo” benchmark contract with no (additional) screening or monitoring. As with the observed contracts in our data, individuals are offered a “uniform” price that varies only by coverage tier, and insurance companies reimburse medical spending, regardless of its origin, based on their contract rules. We estimate that the competitive, average incremental price for the no-deductible plan (relative to the high-deductible plan) is about \$1,570, and that at this competitive price 90 percent of the employees would select the no-deductible plan. We normalize total welfare per employee in this status quo benchmark to be zero, so that we can more easily compare the welfare gains from alternative contract designs.

The second row presents our “perfect screening” counterfactual, which eliminates adverse selection. Specifically, we assume that insurers can observe and price on all the determinants of health care utilization that the individual knows at the time of his insurance choice—i.e., all of the components of $F(\lambda)$ as well as ω . We solve for the incremental price of the no-deductible plan that breaks even for each

³¹ We normalize the price of the lower coverage option to 0. Given our assumptions of CARA utility and a realized utility that is additively separable in income, the price level does not affect plan choice or welfare.

TABLE 10—SPENDING AND WELFARE EFFECTS OF ASYMMETRIC INFORMATION

	Average equilibrium (incremental) premium	No-deductible plan share	Expected spending per employee	Total welfare per employee
(1) “Status quo”: no screening or monitoring	1,568	0.90	5,318	Normalized to 0
(2) “Perfect screening”: premiums depend on $F(\lambda)$ and ω	1,491	0.91	5,248	52
(3) “Imperfect screening”: premiums depend on ω (but not on $F(\lambda)$)	1,523	0.88	5,265	34
(4) “Perfect monitoring”: contracts reimburse only “ λ -related” spending	1,139	0.94	4,185	490
(5) “Imperfect monitoring”: perfect monitoring assumed for choice (but not for utilization)	1,139	0.94	5,327	25

Notes: The table reports the spending and welfare effects from a set of counterfactual contracts described in the text. All exercises are applied to a setting in which the only two options available are the no-deductible plan and the high-deductible plan under the new benefit options (i.e., option 5 and option 1, respectively; see Table 2). Equilibrium premiums are computed as the incremental (relative) premium for the no-deductible plan that equals the expected incremental costs associated with providing the no-deductible plan to those who choose it. The no-deductible plan share is calculated based on the choice probabilities as a function of equilibrium premiums. Expected spending and total welfare are computed based on these choices. Row 1 assumes the “status quo” asymmetric information contracts, with a “uniform” price that varies only by coverage tier. Row 2 assumes “perfect screening,” so that contracts are priced based on ω_i and all components of $F_i(\lambda)$, and adverse selection is eliminated. Row 3 assumes “imperfect screening,” in which contracts are priced based only on ω_i . Row 4 assumes “perfect monitoring” so that moral hazard is eliminated. Specifically, we assume the insurance provider can counterfactually observe (and not reimburse) spending that is associated with moral hazard; spending associated with health—realization of λ —is reimbursed according to the observed contracts. Row 5 assumes “imperfect monitoring” in which, ex ante, individuals choose contracts under the assumption that there will be perfect monitoring (i.e., spending associated with moral hazard will not be reimbursed), but ex post (after they choose their contract but before they make their spending decision) the contracts are changed to be the standard contracts that reimburse all medical spending regardless of its origin.

employee individually, thereby eliminating the adverse selection that arises from uniform pricing. The results indicate that, as expected, the elimination of adverse selection leads to a lower (average) incremental premium for the no-deductible plan, increased coverage (i.e., greater fraction choosing the no-deductible plan), and higher welfare. It also leads to lower expected spending since the risk-based pricing disproportionately shifts higher moral hazard (ω) individuals into lower coverage. We estimate the welfare gain per employee from eliminating adverse selection to be about \$52.³²

Of particular interest is the contribution of eliminating selection on moral hazard—i.e., one component of adverse selection—to the total welfare gain from eliminating adverse selection. Row 3 explores this by reporting the welfare gain from eliminating only selection on moral hazard (ω) but continuing to allow

³²By way of perspective, we calculate the total surplus from perfect screening relative to everyone being in the high-deductible plan to be \$1,084, so that mispricing due to adverse selection appears to reduce welfare by only a small amount relative to the total surplus at stake. Although not the focus of our paper, this finding is consistent with other recent empirical papers on the welfare costs arising from inefficient pricing due to adverse selection; see Einav, Finkelstein, and Levin (2010) for a discussion of some of this recent literature.

selection on health risk ($F(\lambda)$). Specifically, we allow insurers to observe ω and price on it, but not on $F(\lambda)$; thus we are eliminating adverse selection on the slope of health spending with respect to the insurance contract (ω) but not on the level of health spending $F(\lambda)$. This is, of course, not a very sensible scenario, since presumably if insurers could observe ω they could also refuse to reimburse on it, and thus eliminate moral hazard entirely (not just selection on moral hazard). But it is a conceptually useful way to examine the welfare cost of different sources of adverse selection. The results in row 3 suggest that the welfare cost of selection on moral hazard is \$34, or about 65 percent of the \$52 total welfare cost of selection from row 2.

In an analogous fashion, we can investigate the contribution of eliminating selection on moral hazard to the total welfare gain from eliminating moral hazard. In our setting, the welfare gain from eliminating moral hazard stems from two sources: removing the allocative inefficiency that arises from selection on moral hazard and eliminating the traditional moral hazard distortion that comes through socially inefficient overutilization of health care. We show the results from eliminating moral hazard in the fourth row, which presents our “perfect monitoring” counterfactual. Here we assume that insurance coverage only applies to “ λ -related” spending, which in the context of our model means that instead of reimbursing based on actual spending (i.e., reimbursing $m - c_j(m)$), the contracts reimburse $\max\{\lambda, 0\} - c_j(\max\{\lambda, 0\})$ regardless of what the actual spending is. In such situations, optimizing individuals would spend $\max\{\lambda, 0\}$, which would be the socially efficient level of spending. Row 4 of Table 10 indicates that, relative to the status quo (row 1), this elimination of moral hazard reduces spending by more than \$1,100 per employee (column 3) and increases welfare by about \$490 per employee, which is an order of magnitude greater relative to the welfare gain associated with eliminating adverse selection through perfect screening (row 2).

To examine the relative contribution of selection on moral hazard to this welfare cost, in row 5 we again consider an artificial counterfactual. Specifically, we assume that individuals make their contract choices in the first period as if they are faced with the “perfect monitoring” contracts (row 4), but then in the second period make their spending decisions faced with the observed contracts that reimburse in the same manner as the actual contracts (i.e., reimburse based on m rather than based on λ). This allows us to isolate the welfare gain from eliminating solely selection on moral hazard, while preserving the distortion in second-period consumption caused by moral hazard. The results suggest that eliminating selection on moral hazard can achieve welfare gains of \$25 per employee, or only about 5 percent of the total welfare cost of moral hazard (row 4).

Overall, these results suggest that, in our setting, selection on moral hazard contributes nontrivially to the total welfare cost of selection, but contributes much less relative to the total welfare cost of moral hazard. At a broad level, our findings suggest that in thinking about contract design, traditional approaches to combatting moral hazard may well aid in combatting selection, and possibly vice versa. Of course, our quantitative estimates undoubtedly depend on our specific setting (contracts and population) and on our modeling assumptions. While there is not much we can do about the former (at least in the current paper), we investigate the latter in online Appendix F.

VII. Conclusions

This paper takes a first step toward marrying empirical analysis of selection with that of moral hazard. The active (and growing) empirical literature on insurance demand has focused almost exclusively on selection on the level of risk and on risk preferences, and largely abstracted from moral hazard.³³ The large and venerable literature on moral hazard in insurance has largely focused on average moral hazard effects, abstracting from potential heterogeneity as well as potential selection on that heterogeneity. This paper attempts to fill this gap by introducing the possibility that individuals' anticipated behavioral response to insurance affects their contract choice. We suggest that this component of adverse selection can have interesting implications for standard analyses of both selection and moral hazard.

We explored the existence, nature, and implications of selection on moral hazard empirically in the context of employer-provided health insurance in the United States. We estimate substantial heterogeneity in moral hazard and selection on it in our setting, with individuals who have a greater behavioral response to the contract demanding more coverage. We estimate that heterogeneity in moral hazard is roughly as important as heterogeneity in expected health risk in determining whether to buy a higher- or lower-deductible plan. In other words, adverse selection based on the slope of spending (i.e., the incremental spending due to insurance) appears about as quantitatively important in our setting as adverse selection based on the expected level of spending. We illustrate some potential implications of such selection on moral hazard. For example, we estimate that if we ignored selection on moral hazard, we could estimate a spending reduction associated with introducing a high-deductible plan that is substantially larger than what we estimate when we account for the fact that those who select the high-deductible plan have a disproportionately low behavioral response to such cost sharing.

Needless to say, our quantitative estimates are highly specific to our particular population and our particular counterfactual analyses. Nonetheless, at a broad level they illustrate the potential importance of selection on moral hazard for the consequences of both selection and moral hazard and of policies designed to affect these phenomena. They suggest, for example, that efforts to reduce health spending by introducing health insurance options with high consumer cost sharing—such as the high-deductible plans available through Health Savings Accounts—may produce substantially smaller spending reductions than would have been expected based on existing estimates of (average) moral hazard in health insurance. They also suggest that improvements in monitoring technology—traditionally thought of as a way to reduce moral hazard—may have the ancillary benefit of ameliorating some of the efficiency costs of adverse selection.

Given the importance of the topic, we hope that future work will explore selection on moral hazard in other contexts and in other ways. As noted, we know of very little work that even examines heterogeneity in moral hazard, let alone selection of insurance on this heterogeneity. Both the approaches taken in this paper and those suggested (but not explored) by Einav, Finkelstein, and Cullen (2010) for estimating

³³ See Einav, Finkelstein, and Levin (2010) for a recent discussion of this literature.

heterogeneity in moral hazard and its correlation with demand should be fruitful to apply in other settings. In addition, our analysis has focused exclusively on some of the implications of selection on moral hazard for a given set of contracts; it would be interesting to consider, both theoretically and empirically, the implications of selection on moral hazard for richer analyses of contract design.

REFERENCES

- Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark R. Cullen.** 2012. "Moral Hazard in Health Insurance: How Important is Forward Looking Behavior?" National Bureau of Economic Research Working Paper 17802.
- Arrow, Kenneth.** 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53 (5): 941–73.
- Arrow, Kenneth.** 1971. *Essays in the Theory of Risk-Bearing*. Amsterdam: North-Holland Publishing Company.
- Bajari, Patrick, Han Hong, Ahmed Khwaja, and Christina Marsh.** 2010. "Moral Hazard, Adverse Selection and Health Expenditures: A Semiparametric Analysis." <http://www.econ.umn.edu/~bajari/research/hazard.pdf> (accessed November 7, 2012).
- Berry, Steve, and Ariel Pakes.** 2007. "The Pure Characteristics Model of Demand." *International Economic Review* 48 (4): 1193–1225.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney.** 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review* 102 (7): 3214–48.
- Cardon, James H., and Igal Hendel.** 2001. "Asymmetric Information in Health Insurance: Evidence from The National Medical Expenditure Survey." *RAND Journal of Economics* 32 (3): 408–27.
- Carneiro, Pedro, James J. Heckman, and Edward Vytlačil.** 2011. "Estimating Marginal Returns to Education." *American Economic Review* 101 (6): 2754–81.
- Carlin, Caroline, and Robert Town.** 2010. "Adverse Selection: The Dog That Didn't Bite." Unpublished.
- Chandra, Amitabh, Jonathan Gruber, and Robin McKnight.** 2010. "Patient Cost-Sharing and Hospitalization Offsets in the Elderly." *American Economic Review* 100 (1): 193–213.
- Chiappori, Pierre-André, and Bernard Salanié.** 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy* 108 (1): 56–78.
- Cohen, Alma, and Liran Einav.** 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review* 97 (3): 745–88.
- Cutler, David, and Richard Zeckhauser.** 2000. "The Anatomy of Health Insurance." In *Handbook of Health Economics*. Vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 563–644. Amsterdam: North-Holland.
- Dubin, Jeffrey A., and Daniel L. McFadden.** 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52 (2): 345–62.
- Ehrlich, Isaac, and Gary S. Becker.** 1972. "Market Insurance, Self-Insurance, and Self-Protection." *Journal of Political Economy* 80 (4): 623–48.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010. "Estimating Welfare in Insurance Markets using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877–921.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin.** 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics* 2 (1): 311–36.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf.** 2010. "Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from The U.K. Annuity Market." *Econometrica* 78 (3): 1031–92.
- Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark R. Cullen.** 2011. "Selection on Moral Hazard in Health Insurance." National Bureau of Economic Research Working Paper 16969.
- Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen.** 2013. "Selection on Moral Hazard in Health Insurance: Dataset." *American Economic Review*. <http://dx.doi.org/10.1257/aer.103.1.178>.
- Fang, Hanming, Michael Keane, and Dan Silverman.** 2008. "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market." *Journal of Political Economy* 116 (2): 303–50.
- Finkelstein, Amy, and Kathleen McGarry.** 2006. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." *American Economic Review* 96 (4): 938–58.

- Goettler, Ronald L., and Karen Clay.** 2011. "Tariff Choice with Consumer Learning and Switching Costs." *Journal of Marketing Research* 48 (4): 633–52.
- Handel, Benjamin.** 2012. "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts." http://emlab.berkeley.edu/~bhandel/wp/Handel_ASIN_82012.pdf (accessed November 7, 2012).
- Heckman, James J., Sergio Urzua, and Edward Vytlacil.** 2006. "Understanding Instrumental Variables in Models with Essential Heterogeneity." *Review of Economics and Statistics* 88 (3): 389–432.
- Holmström, Bengt.** 1979. "Moral Hazard and Observability." *Bell Journal of Economics* 10 (1): 74–91.
- Keane, Michael, and Robert Moffitt.** 1998. "A Structural Model of Multiple Welfare Program Participation and Labor Supply." *International Economic Review* 39 (3): 553–89.
- Keeler, Emmett B., and John E. Rolph.** 1988. "The Demand for Episodes of Treatment in the Health Insurance Experiment." *Journal of Health Economics* 7 (4): 337–67.
- Manning, Willard, Joseph Newhouse, Naihua Duan, Emmett Keeler, Arleen Leibowitz, and Susan Marquis.** 1987. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review* 77 (3): 251–77.
- Miravete, Eugenio J.** 2003. "Choosing the Wrong Calling Plan? Ignorance and Learning." *American Economic Review* 93 (1): 297–310.
- Newhouse, Joseph.** 1993. *Free for All? Lessons from the RAND Health Insurance Experiment*. Cambridge, MA: Harvard University Press.
- Pauly, Mark.** 1968. "The Economics of Moral Hazard: Comment." *American Economic Review* 58 (3): 531–37.
- Wooldridge, Jeffrey M.** 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.