RAND Journal of Economics
Vol. 45, No. 4, Winter 2014
pp. 747–763

Moral hazard, adverse selection, and health expenditures: A semiparametric analysis

Patrick Bajari*

Christina Dalton**

Han Hong***

and

Ahmed Khwaja****

Theoretical models predict asymmetric information in health insurance markets may generate inefficient outcomes due to adverse selection and moral hazard. However, previous empirical research has found it difficult to disentangle adverse selection from moral hazard in health care consumption. We propose a two-step semiparametric estimation strategy to identify and estimate a canonical model of asymmetric information in health care markets. With this method, we can estimate a structural model of demand for health care. We illustrate this method using a claims-level data set with confidential information from a large self-insured employer. We find significant evidence of moral hazard and adverse selection.

1. Introduction

■ Despite a large theoretical literature predicting asymmetric information leads to inefficiencies in insurance markets (Akerlof, 1970; Arrow, 1963; Spence and Zeckhauser, 1971), it is empirically difficult to distinguish between adverse selection and moral hazard. In the context of health insurance, Cutler and Zeckhauser (2000) review an extensive literature that finds evidence of adverse selection based on the positive correlation between insurance contract generosity and adverse outcomes, and moral hazard based on coinsurance elasticity of the demand for medical care. However, Chiappori and Salanié (2003) show that, under moral hazard, contract generosity leads to adverse risk outcomes, whereas under adverse selection, the causality is reversed. This

* University of Washington and NBER; bajari@uw.edu.
** Wake Forest University; marshcl@wfu.edu.
*** Stanford University; hanhong@stanford.edu.
**** Yale University; ahmed.khwaja@yale.edu.
We have benefited from the comments of participants at various conferences and seminars. We acknowledge excellent research assistance from Ivan Shaliastovich and, especially, Alvin Murphy. Bajari and Hong would like to thank the National Science Foundation for generous research support with NSF Research Grants no. SES-0721015 and no. SES-1024504. We thank three anonymous referees and the editor for their helpful suggestions.

leads to observational equivalence between the two hypotheses. Predictions of insurance theory and efficient insurance market regulation depend on whether adverse selection or moral hazard is more prevalent. In this article, we introduce an analytical method that can separately identify these two effects.

Our approach incorporates methods from the auctions literature (e.g., Campo et al., 2011) to offer a new way to estimate preferences for medical utilization choices and the underlying unobserved health distribution of an insured population. The current literature is almost exclusively based on modelling selection into plans and consequent health care utilization using a two-stage estimation approach. See, for example, Cardon and Hendel (2001), Einav, Jenkins, and Levin (2012).[1] Our method complements this literature by providing a way to estimate preferences for health care utilization, which can be used to quantify the extent of moral hazard, and in turn assess the presence of adverse selection. An advantage of our identification is that recovered preferences and the degree of selection are based on *ex post* realized health shocks in contrast to *ex ante* information, without having to recover health plan choice preferences.

We begin with a model of consumers' health care choices given plan contract characteristics. A consumer has preferences over both health care expenditure and aggregate consumption, which are influenced by the consumer's latent health status. The consumer's budget constraint is specific to her insurance plan and also incorporates uncertainty in reimbursement for health expenses. Adverse selection arises because consumers have private information about their latent health status, which is unobserved by the insurer. Moral hazard results because consumers do not pay the full cost of the health care coverage. We then estimate risk parameters semiparametrically and use these parameters to recover the latent health status distribution. We use the estimated plan-specific latent health distributions to test for adverse selection across plans. Last, given the estimated health status distribution, utility parameters, and observed consumption choices, we are able to infer the level of moral hazard with a counterfactual replicating a social planner's optimal consumption allocation.

We demonstrate this method by estimating each step on a confidential claims-level data set from a large self-insured employer. The data contains a high level of detail on spending and reimbursement, as well as individual demographic and health status indicators. Our results indicate both significant adverse selection and moral hazard. In a nonparametric test for adverse selection, we find that healthy individuals sort into a generous, but restrictive plan. In our plans, median levels of moral hazard are over 40% of health expenditures.

Our work contributes to the analysis of asymmetric information in insurance markets in three significant ways. First, we allow for both adverse selection and moral hazard due to asymmetric information in estimating the model. Second, our estimation strategy is semiparametric, in contrast to earlier work relying on parsimoniously specified parametric models. This is important because theory provides little guidance about which parametric distributions for latent health shocks are *a priori* most plausible. Semiparametric estimation also allows us to specify plans' reimbursement schedules flexibly, allowing for nonlinearities caused by common characteristics such as deductibles and copays. These nonlinearities are not captured by the more restrictive specifications of previous work. Third, our method only necessitates a single, relatively weak identification assumption. Previous structural approaches rely on strong identifying assumptions. In this article, the sole identifying assumption is that the distribution of health shocks is invariant over a large population during short consecutive time periods.

Our research is novel in that it develops a tractable estimation procedure under minimal parametric assumptions to simultaneously examine adverse selection and moral hazard in health insurance contracts. It provides an important framework for similar analysis in other contexts, especially with cross-section data, where distortions exist due to asymmetric information. The rest of the article is organized as follows. The model is discussed in Section 2 and identification

[1] Einav et al. (2013) also present a selection plus utilization approach and discuss explicitly the limitations of estimating moral hazard and adverse selection separately.

is outlined in Section 3. Our tests for adverse selection and moral hazard are described in Section 4. Section 5 lays out the estimation steps with our claims data as illustration, and Section 6 briefly describes the results for our data. Section 7 concludes.

2. Model

■ We develop a canonical model of demand for health care to incorporate the essential features of asymmetric information in health insurance markets. This model is general, building on the work of Blomqvist (1997), Cardon and Hendel (2001), and Spence and Zeckhauser (1971), though we present a more specific form for the empirical application. In this model, consumers act directly to maximize consumer utility. Consumers choose the optimal amount of health services, m, and composite good consumption c, subject to a budget constraint.

Our model differs from the standard neoclassical model of consumer choice in two important respects. First, the optimal choice of m and c depends on a consumer's latent health status, which we denote as the scalar θ. We interpret θ as a preference shock for health services. The higher the value of θ, the higher the utility from health services. Asymmetric information results because consumers know their value of θ, but insurers do not observe θ. The second distinct feature is the budget constraint, which incorporates the more complicated pricing framework of insurance compared to that of ordinary consumer goods. Health insurance plans introduce nonlinearities into the budget constraint through features such as deductibles and coverage caps (Keeler et al., 1977).[2] Additionally, the consumer's out-of-pocket expenses may be uncertain at the time of the health care choice. We allow this out-of-pocket expense to be stochastic. This generalizes previous research, which typically assumes that the consumer has known and certain costs of health care.

□ **Consumer choice.** The consumer's utility function is specified as:

$$\mathrm{U(c, m; \theta, \gamma) = F(c, (1 - \theta), \gamma_1) + H(m, \theta, \gamma_2)}, \tag{2.1}$$

where $F(.)$ is utility from composite good consumption good, $H(.)$ is utility from health good consumption, θ is the health status parameter, and γ_1 and γ_2 are risk aversion parameters. Utility is additively separable, and $F(.)$ and $H(.)$ obey standard utility function properties of diminishing marginal returns in both goods. The latent health status parameter θ does not have a utility value in itself but lies between [0,1] and indexes the weight that the consumer places on consumption of health services versus the nonhealth commodity good. If θ is close to one, the consumer values the health good, m, more and the $(1 - \theta)$ term in composite consumption becomes small. The utility function that we specify depends directly on health expenditure and the consumption of the composite good. It implies an indirect utility function for each patient when each patient maximizes utility subject to the budget constraint implied by the price and income levels.

The parameters γ_1 and γ_2 enable multidimensional risk aversion, describing the consumer's risk attitude with respect to aggregate consumption and health, respectively. Separate risk parameters are important because out-of-pocket expenditures associated with m may be uncertain. We assume constant relative risk aversion (CRRA) in c and m.

Consumers are heterogeneous because θ varies across individuals. Denote the probability density function of the health status distribution as $g(\theta)$. In principle, we could let $g(\theta)$ depend on observed covariates associated with health status, such as education level or income. However, in the absence of a theoretical foundation for the choice of these covariates, any ad hoc specification involves the risk of misspecification. Our estimation will reveal the general shape of this distribution, and, as a result, we do not need to specify any particular assumptions between certain covariates and the θ distribution.

[2] See for example, Bajari et al. (2013) and Dalton (2014) for a different approach based on a regression discontinuity design to estimate price sensitivity under nonlinear contracts.

□ **Budget constraint.** The budget constraint specifies that the consumer's total expenditures on the composite good plus health good must be less than her income, y, after deducting insurance premiums from plan j, p_j.[3]

$$\mathrm{c} + \mathrm{m}(1 - \mathrm{a_j}) \leq \mathrm{y} - \mathrm{p_j} \tag{2.2}$$

The consumer's out-of-pocket cost for health care goods is determined by the reimbursement rate of insurance plan j, which we label as the scalar a_j. If a consumer chooses health good consumption m, then the insurer will cover $a_j m$ of these expenses. As a result, the consumer must pay for $m(1 - a_j)$ out-of-pocket.

□ **Incorporating uncertainty.** The health care consumer faces the difficulty that reimbursement a_j may be uncertain when choosing m. Medical plans are long and typically quite complicated documents written by insurers, their executives, and attorneys. It is unlikely that a typical consumer invests the resources to understand what these medical plans cover in all states of the world. Furthermore, health providers and insurance administrators often negotiate the final reimbursement after services were consumed.

As we are ultimately interested in consumer demand, it is important to model reimbursement, a_j, from the perspective of the consumer. Therefore, we model a_j as a conditional probability $f_j(a_j|m)$. This allows the generosity of benefits to depend on the plan j chosen by the consumer, but also to depend on the choice of health consumption, m. For example, many plans require a consumer to pay a fixed fee for a doctor's visit. Plans may display features such as deductibles or coverage gaps, which do not reimburse within a certain range of expenditures. Our framework accommodates these complications by allowing the reimbursement rate to depend on m.

We will use the observed distribution of reimbursements and flexible, nonparametric methods to identify $f_j(a_j|m)$. We allow the conditional distribution to be data driven and consistent with the standard economic assumption that consumers have rational expectations. Consumer beliefs about a_j must be consistent with the observed outcomes.[4]

The consumer makes her choice under uncertainty. The consumer first draws a value of θ from $g(\theta)$, then makes her choice of m. The reimbursement rate a_j is realized from the distribution $f_j(a_j|m)$. As preferences are strictly increasing, the budget constraint binds and c is determined by the equation $c = y - p_j - m(1 - a_j)$.

Let $EU(m, p_j, y; \theta, \gamma)$ denote the consumer's expected utility when choosing m. Substituting in for c yields the following:

$$\mathrm{EU(m, p_j, y; \theta, \gamma)} = \int \mathrm{F((y - p_j - m(1 - a_j)), (1 - \theta), \gamma_1) f_j(a_j|m) da_j} + \mathrm{H(m, \theta, \gamma_2)}. \tag{2.3}$$

In computing expected utility, the consumer integrates over reimbursement a_j using the distribution $f_j(a_j|m)$. If the realization of a_j is closer to one, the value of composite good

[3] The utility function here is conditional on an individual's plan choice. We do not model the plan choice because the optimal choice of health expenditure, m, conditional on an individual's plan choice, will be sufficient to identify our model parameters. As such, our model does not fully capture, and the estimated preferences do not reflect, risk aversion at the insurance plan choice stage. Our approach has other advantages, however. In particular, this approach allows us to estimate the unobserved health shock distribution without needing to model the choice of insurance. The role of uncertainty and risk aversion in insurance choice is an important and interesting issue by itself. Although it is well beyond the scope of the current model and approach, future work could investigate the possibility of embedding our model in a richer framework that includes the choice of a health plan at an earlier stage along the lines of, for example, Cardon and Hendel (2001), Khwaja (2010), Bundorf, Levin, and Mahoney (2012), and Handel (2013), who model health plan choices in the presence of asymmetric information and medical expenditure risk.

[4] As an alternative, it might be interesting to allow consumers to have biased or irrational beliefs about the determination of a_j. However, our model is flexible and comes close to exhausting the degrees of freedom in the data. The identification of such irrational beliefs would therefore be tenuous. There is no consensus about a theoretical framework which would provide a plausible basis for the *a priori* specification of how consumers bias their beliefs in the context of health care. As a consequence, we use the more common assumption that consumers have beliefs that are consistent with *ex post* outcomes.

consumption will be larger, all else held equal. The value of expected utility depends on the consumer's attitude toward risk. For example, the more risk adverse the consumer is toward uncertainty in composite good consumption, the less health good she would consume. Utility also depends on income and premiums through $y - p_j$. For example, households with lower incomes are more adversely impacted by a low realization of a_j.

3. Identification

■ Our next step is to use the framework above to estimate the health status distribution.

□ **Inferring θ from observed choices.** The first step to constructing the latent θ distribution is to solve for first-order conditions of the consumer's expected utility with respect to the choice of the health good, as described in equation (2.3). After we set the first-order conditions equal to zero, we can rearrange the resulting terms so the parameter θ is a function, I, of known data components and unknown parameters:

$$\theta = \mathrm{I}(\mathrm{y}, \mathrm{p_j}, \mathrm{m}, \mathrm{a_j}, \mathrm{f_j}(\mathrm{a_j}|\mathrm{m}), \frac{\partial \mathrm{f_j}(\mathrm{a_j}|\mathrm{m})}{\partial \mathrm{m}}; \gamma_1, \gamma_2). \tag{3.1}$$

The left-hand side of equation (3.1) is the consumer's private health status. The observables in the right-hand side of equation (3.1) are income, y, premiums, p_j, the individual consumer's health spending, m, and the individual consumer's reimbursement of health spending, a_j. The conditional reimbursement probability distribution, $f_j(a_j|m)$, can also be directly observed in claims data. Finally, the derivative of the reimbursement probability with respect to health spending, $\frac{\partial f_j(a_j|m)}{\partial m}$, can be approximated from $f_j(a_j|m)$. The unknowns that remain in equation (3.1) are the risk parameters, γ_1 and γ_2.

The rationale underlying our estimator is similar to that used in auction models, especially Campo et al. (2011) and Guerre, Perrigne, and Vuong (2000). In auction models, estimation centers around finding consumers' unobserved valuations of an item for bid; here, we are estimating the consumer's unobserved health status that leads to health care purchases. Under suitable regularity conditions, for a fixed value of γ, we will always be able to find a value of θ that rationalizes the consumer's choice. As a result, it follows that the assumption of utility maximization alone will not be adequate to identify both the health status distribution $g(\theta)$ and risk parameters γ. Our approach to identification is to impose additional moment restrictions.

□ **Identifying assumption.** Our identification strategy is similar to the risk averse auction model in Campo et al. (2011). We approach identification by imposing an additional moment restriction, in particular that $g(\theta)$ does not depend on time. Given the assumption of a stable health status distribution, we can then use this distribution's moments across years to identify the only remaining unknowns, the risk parameters γ. In nonlinear parametric models, global identification is generally difficult to verify, and Campo et al. (2011) also make use of a set of nonlinear moment conditions to estimate risk aversion parameters.[5]

Intuitively, this assumption that $g(\theta)$ doesn't depend on time requires that the health status distribution calculated over a large population does not change over short consecutive time periods. For example, our application uses a three-year panel on incomes and health care choices from a large Minnesota-based employer. In this application, we argue this restriction is reasonable

[5] Campo et al. (2011)'s "high-level" parametric identifying condition A1 (iv) implies the order condition for identification, but does not explicitly provide a rank condition. For a special case of the constant relative risk aversion (CRRA) utility function, their moment conditions become linear so that the rank condition can be directly tested using the observed data. However, Campo et al. (2011) candidly acknowledge that "considering another parametric specification...would lead to a nonlinear system of equations in (the risk aversion parameters) for which local identification conditions can be obtained through the usual 'rank' conditions," without providing explicit details for the verification of the rank conditions.

because the population of employees is very large and there is no reason to expect large fluctuations in the health status distribution as reflected by $g(\theta)$ among this group within a reasonably short time period.[6]

Although overall levels of health status remain stable, identifying variation exists in individual health shocks. Premiums, p_j, and reimbursement rates, $f_j(a_j|m)$, also vary from year to year, and variation exists in individual income through promotions or job change within the organization.[7]

Intuitively, unobserved health status θ is the analogue of the unobserved error term in a structural linear demand equation. The effective instruments, which are the year dummies, are associated with exogenous shifts in the distribution of income and the relative price of medical care that are uncorrelated with health status. Thus, we can hold the health status "fixed" although the relative prices are shifted by the instruments, which leads to co-movement in medical expenditures and consumption. Such co-movements allow us to identify the utility parameters. Once we have estimated the utility parameters, the distribution of health status can be backed out from the first-order condition. The difference between our method and the conventional two-stage least squares estimator is that, instead of relying on a reduced-form specification of a linear functional form for the demand equation, we use the optimality condition for a risk-averse consumer to derive the functional form of the demand equation.

4. Moral hazard and adverse selection

■ After estimating the last remaining unknowns, the risk parameters γ, we can now construct the estimated health shock distribution $g(\theta)$, substituting each individual's data into an estimated version of equation (3.1) to obtain individual θs. The inability of insurers to observe these health status parameters leads to the two primary issues in insurance, adverse selection and moral hazard. We can now use our estimated θ distribution to answer questions about these two phenomena.

□ **Adverse selection.** Adverse selection occurs when insurers are not able to distinguish between (or restrict the choices of) consumers with different values of θ. If an insurer designs a plan j with generous features, such as high reimbursement or better access to care, these features attract consumers with high values of θ, even if plan j was designed for consumers with low values of θ. With our estimated latent θ distribution, we can test for sorting of θ types between different employer plans.

We propose a distribution-free test for adverse selection. The presence of adverse selection causes consumers to sort across different plans based on their latent health status (e.g., Rothschild and Stiglitz, 1976). In our framework, this implies that the distribution of the latent health status variable varies across health plans. Hypothesis testing reveals these patterns. The null hypothesis is that the estimated latent health distribution within one plan looks similar to the distribution of the other plan. The alternative hypothesis is that the distributions were actually drawn from different populations.

To test the direction of sorting, we can test the distributions for stochastic dominance using hypothesis testing. A more severe health status enters utility as a larger value of θ. Thus, the cdf of θs of a plan with better health statuses (lower values of θ) should stochastically dominate the cdf of a plan with worse health statuses (greater values of θ).

[6] We researched any possible epidemics and significant changes, such as innovations in medical technology, during this time period. We found no evidence of either type that would lead to a large shift in the health distribution in this population during the three-year period.

[7] In theory, another source of variation in the latent health distribution could be the option of not insuring at all. However, our application is a large self-insured employer, so selection effects are limited because most employees are covered by the plans offered by the employer. In our application, we use those plans which capture most of the employee base, insuring over 80% of employees.

□ **Moral hazard.** Moral hazard occurs in this modelling framework because consumers pay only a proportion a_j of each unit of health care m, not the full cost incurred by the insurer. In our model, insurers establish premiums, p_j, and reimbursement probabilities, $f(a_j|m)$, but, once consumers have chosen plan j, the insurer cannot contract the amount of care, m, a consumer chooses within plan j. Moral hazard measures how consumers' choices change once insurance insulates them from the true cost of care. This definition is commonly used in the health economics literature, as established by Pauly (1968).[8]

Our test of moral hazard is as follows. We construct a counterfactual scenario where the relative prices of health care and the composite good are each set to one, so an additional dollar of income buys the same amount of each good. The consumer is given a lump-sum income transfer equal to the observed amount of original health care consumption. This lump-sum transfer is denoted as $T = a_j \cdot m_1$, where m_1 is the original health care consumption. This ensures the original consumption bundle, (m_1, c_1), remains affordable to the consumer.

The counterfactual budget constraint is then:

$$c_2 + m_2 \leq y - p_j + T,$$

where T represents the lump-sum transfer to an individual consumer. The variables c_2 and m_2 are the composite good consumption and health spending choices in the counterfactual environment. The consumer no longer faces reimbursement schedule a_j but receives a lump-sum transfer of T and pays for all care out-of-pocket.

Given the estimated health shocks, $\hat{\theta}$ and utility parameters $\hat{\gamma}$, we allow the consumer to reoptimize with respect to her choice of health care consumption. The extent of moral hazard or "overconsumption" is captured by the subsequent change in the optimal consumption bundle–the difference between the original choice, m_1, and the new choice, m_2.

In the counterfactual scenario, consumers are no longer restricted to spend their reimbursement income, T, on health expenditure. If the consumer chooses to spend part of the unrestricted lump-sum transfer, T, on consumption of the aggregate good, this reveals how purchasing health care through insurance influences behavior.[9]

5. Empirical application

■ To illustrate how to estimate the $g(\theta)$ distribution and then use these estimated values to answer questions about adverse selection and moral hazard, we apply our method to a data set of health insurance claims.

□ **Data.** We use a detailed, confidential claims-level data set from a large self-insured employer for the years 2002–2004.[10] The employer has offices in several locations and the complete claims data covers over 39,000 total beneficiaries. Each claim includes total approved health care spending, total reimbursement, and cost variables such as the coinsurance or copayment paid, and the amount of deductible used. Costs and spending amounts are aggregated for each beneficiary over the year for a precise measure of total spending by the employer and total out-of-pocket cost to the beneficiary. Employee information includes age and salary. Characteristics such as premiums, employee contribution, and cost structure come directly from the employer's enrollment materials.

[8] This definition of moral hazard is commonly referred to as *ex post* moral hazard, in contrast to *ex ante* moral hazard. *Ex post* moral hazard focuses on consumer behavior once an insurance-related incident has occurred. Other fields, such as contract theory, often focus on *ex ante* moral hazard, which refers to "hidden actions" of consumers before the incident occurs.

[9] This approach incorporates the insights of sequential contracting, where an agent selects her best action within a previously chosen contract (see Courty and Li, 2000; Dai, Lewis, and Lopomo, 2006).

[10] We thank Robert Town and Caroline Carlin for their assistance in accessing this data. For more detail on this data set, refer to Carlin and Town (2010).

TABLE 1 Total Enrollees By Plan

Plan	2002	2003	2004	Total
HMO	4032	4126	3845	12,003
PPO1	483	469	388	1340
PPO2	523	539	230	1292
Total	5038	5134	4463	14,635

TABLE 2 Summary Statistics By Plan

Plan	Variable	Mean	Standard Deviation
HMO	Total health spending	$2845	$7459
	Total reimbursement	$2659	$7139
	Age	41.7	12.0
	Health status proxy	1.4	2.8
PPO1	Total health spending	$3964	$7559
	Total reimbursement	$3657	$7555
	Age	46.4	11.5
	Health status proxy	2.2	3.2
PPO2	Total health spending	$4700	$9243
	Total reimbursement	$4449	$9151
	Age	47.0	10.7
	Health status proxy	2.5	3.8

Observations: HMO=12,003, PPO1=1,340, PPO2=1,292.

We focus on the three most populous plans which comprise over 80% of enrollment. These plans are a Health Maintenance Organization (HMO) plan and two tiered Preferred Provider Organization (PPO) plans, referred to as "PPO1" and "PPO2." We use only employees enrolled under single coverage with continuous enrollment for the entire year.[11] Table 1 displays enrollment statistics for the sample. Over the three-year period, the data set has over 14,000 single-coverage enrollees. The HMO plan has the largest enrollment, with over 80% of enrollees each year. Enrollments in the PPO1 and PPO2 plans range between 230 and 539.

Table 2 displays summary statistics for each of the three plans for total health spending, total reimbursement, age, and a health status proxy of weighted diagnosis codes. Increasing values of the health status proxy indicate increasing severity.[12] Average health care spending is less than $5000 in all three of the plans, but the tiered PPO plans have higher average spending than the HMO plan. For each plan, total reimbursement is slightly less than total health spending, which reflects enrollees' out-of-pocket costs.

Comparing demographic variables across the plans, there is preliminary evidence of adverse selection. The HMO plan's enrollees are younger compared to PPO1 and PPO2. The average of the health status proxy variable is also much lower for the HMO plan, with an average health status proxy value of 1.4. The two PPO plans' average health status proxy values are both well over 2. Although these differences in demographic characteristics do not prove adverse selection on unobservable illness levels, they offer evidence that plans do vary in observable characteristics.

□ **Estimation.** We use a two-step semiparametric estimation strategy to recover the underlying parameters of the consumer's utility function and the distribution of latent health status. An individual observation is a full-year single coverage enrollee i, in a plan j, for a given year t.

[11] Continuous single coverage replicates our specified model, but more consumer categories could be analyzed with simple modifications to our framework.

[12] The health proxy variable was created for each enrollee using the Johns Hopkins University ACG (Adjusted Clinical Groups) Case-Mix System (v.6), which is widely used in the health care sector to incorporate individual-level diagnostic claims data to predict future medical expenditures accounting for various comorbidities.

Estimate conditional reimbursement distributions. The first estimation step is to calculate the conditional reimbursement distribution, $f_{jt}(a_j|m)$, for each plan and year combination. We nonparametrically estimate these distributions using observational data on the reimbursements, a_j, and the spending levels, m. The *ex post* realized reimbursement, a_j, is calculated for each enrollee as a percentage of total health spending. The nonparametric kernel estimation uses the realized reimbursements to calculate the joint probability of a given reimbursement percentage and a given level of health spending, $\hat{f}_{jt}(a_{jt}, m_{jt})$.[13] We then nonparametrically calculate the health spending distribution, $\hat{f}_{jt}(m_{jt})$. The estimated conditional distribution, $\hat{f}_{jt}(a_{jt}|m_{jt})$ is

$$\hat{\mathrm{f}}_{\mathrm{jt}}(\mathrm{a}_{\mathrm{jt}}|\mathrm{m}_{\mathrm{jt}}) = \frac{\hat{\mathrm{f}}_{\mathrm{jt}}(\mathrm{a}_{\mathrm{jt}}, \mathrm{m}_{\mathrm{jt}})}{\hat{\mathrm{f}}_{\mathrm{jt}}(\mathrm{m}_{\mathrm{jt}})}.$$

The conditional distribution is calculated for a grid of reimbursement percentages and health expenditure categories.[14] Figure 1 displays the conditional distributions $\hat{f}_{jt}(a_{jt}|m_{jt})$ for each of the three plans in 2002. The bottom axes show the percentage reimbursed, increasing from 0 (0%) to 1 (100%) on the left, and the levels of health spending, increasing from \$0 to over \$5000 on the right. The surface rising above these axes shows the conditional probabilities of a certain reimbursement percentage given the health spending level. Across all years, the lowest reimbursement percentages occur at the lowest level of health spending, and vice versa. To compare the plans, notice that the surface for the HMO plan rises higher than the other plans, meaning the HMO's conditional probability of full reimbursement at the highest spending level is nearly at 1, or 100% certainty. The same conditional probability for PPO1 only reaches a maximum value of just over 0.7, or 70%, and PPO2's same probability never reaches above 70%. Finally, we obtain $\frac{\partial \hat{f}_{jt}(a_{jt}|m_{jt})}{\partial m_{jt}}$ by applying an approximate derivative to all grid points of the conditional distribution.[15]

Mapping the unobservable θ to observables. To illustrate our method, we choose a simple functional form for consumer utility that fits the general utility framework presented in Section 2. Utility is

$$\mathrm{U}(\mathrm{c}, \mathrm{m}; \theta, \gamma) = (1-\theta)\frac{\mathrm{c}^{1-\gamma_1}}{1-\gamma_1} + \theta\frac{\mathrm{m}^{1-\gamma_2}}{1-\gamma_2}.$$

We incorporate uncertainty and substitute for composite good consumption when the budget constraint binds. Expected utility is

$$\mathrm{EU}(\mathrm{m}; \mathrm{p_j}, \mathrm{y}, \theta, \gamma) = \int (1-\theta)\frac{(\mathrm{y} - \mathrm{p_j} - \mathrm{m}(1-\mathrm{a_j}))^{1-\gamma_1}}{1-\gamma_1}\mathrm{f_j}(\mathrm{a_j}|\mathrm{m})\mathrm{da_j} + \theta\frac{\mathrm{m}^{1-\gamma_2}}{1-\gamma_2}. \quad (5.1)$$

The consumer maximizes her expected utility with respect to her choice of health care spending, m. We set the derivative of expected utility equal to zero, and rearrange the FOC where the θ term is on the left as a function of all the observables and the risk parameters. We then substitute in any observable variables with their estimated counterpart to obtain the following estimation equation for an individual consumer's $\hat{\theta}_i$ value:

$$\hat{\theta}_{\mathrm{i}} = \frac{\hat{\mathrm{I}}}{\hat{\mathrm{I}} - \frac{(1-\hat{\gamma}_2)}{\mathrm{m}_{\mathrm{ijt}}^{\hat{\gamma}_2}}}, \quad (5.2)$$

[13] The bandwidth is the optimal bandwidth rule of thumb suggested by Bowman and Azzalini (1997).

[14] We display results using a 128 x 128 grid of a_js and m_{jt}s. Both more and fewer categories produced very little change in the resulting estimates.

[15] These estimation results use an adaptive Simpson quadrature for the derivative.

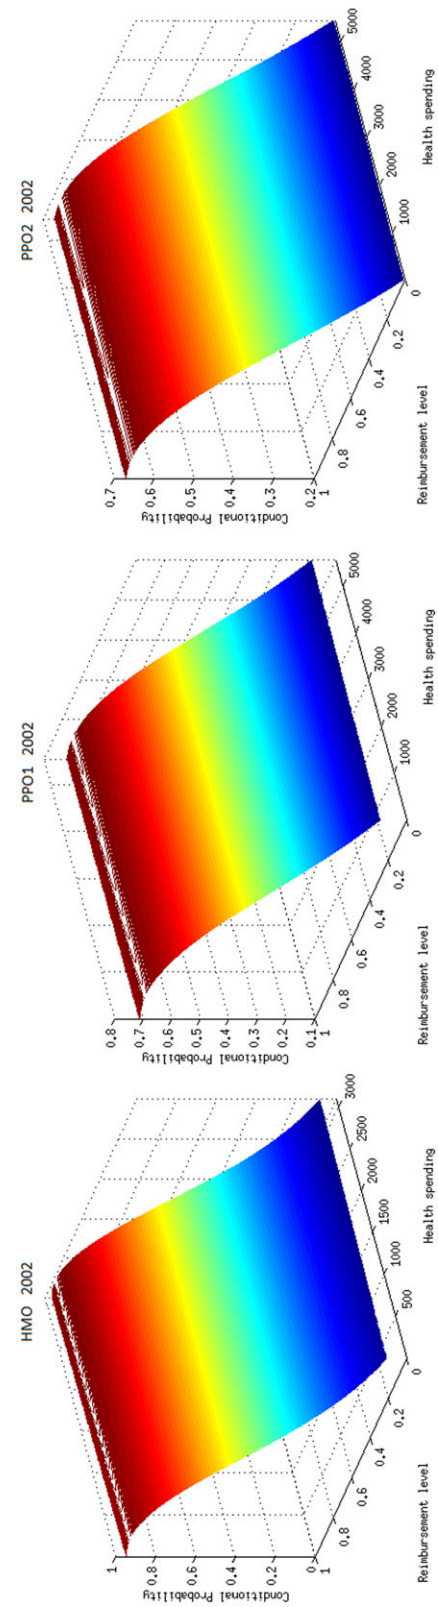

FIGURE 1

Conditional Reimbursement Distributions, 2002

TABLE 3 Health Status Proxy By Year

Year	p25	p50	p75	Mean	Standard Deviation	*N*
2002	0.20	0.91	1.49	1.64	3.12	5038
2003	0.20	0.91	1.49	1.60	2.81	5134
2004	0.20	0.91	1.49	1.57	2.92	4463
Total	0.20	0.91	1.49	1.60	2.96	14,635

TABLE 4 Risk Parameter Estimation Sample

Plan	2002	2003	2004	Total
HMO	2690	2791	2631	8112
PPO1	351	341	295	987
PPO2	392	392	174	967
Total	3433	3533	3100	10,066

where

$$\hat{\mathrm{I}} = \int [-(1-\hat{\gamma_1})(1-\hat{\mathrm{a}}_{\mathrm{jt}})(\mathrm{y}_{\mathrm{it}} - \mathrm{p}_{\mathrm{jt}} - \mathrm{m}_{\mathrm{ijt}}(1-\hat{\mathrm{a}}_{\mathrm{jt}}))^{-\hat{\gamma}_1} \hat{\mathrm{f}}_{\mathrm{jt}}(\mathrm{a}_{\mathrm{jt}}|\mathrm{m}_{\mathrm{jt}})$$

$$+(y_{it} - p_{jt} - m_{ijt}(1-\hat{a}_{jt}))^{1-\hat{\gamma}_1} \frac{\partial \hat{f}_{jt}(a_{jt}|m_{jt})}{\partial m_{jt}}]da_{jt}. \quad (5.3)$$

Data components are health care spending m_{ijt}, income y_{it}, and premiums p_{jt}. Health care spending and income are unique to an individual i in plan j in a given year t. Premiums are the same for all individuals in a plan j in a given year t. The conditional reimbursement distributions $\hat{f}_{jt}(a_{jt}|m_{jt})$ and $\frac{\partial \hat{f}_{jt}(a_{jt}|m_{jt})}{\partial m_{jt}}$ were estimated from the data in the previous step. The remaining unknowns in equation (5.2) are the risk aversion parameters $\hat{\gamma_1}$ and $\hat{\gamma_2}$.

Estimate risk parameters using GMM. For each single-coverage full-year enrollee i in a given plan j and a given year t, we now have an expression for the enrollee's latent health status. We use these health status expressions in conjunction with the identifying assumption to estimate utility parameters using a Generalized Method of Moments (GMM) framework.

Our identifying assumption requires the latent health status distributions to be equal to each other across years. As a check on the identification assumption, Table 3 displays the health status proxy from the data for each year.[16] The population quartiles are identical across all years. Health spending distributions often have long right tails. Despite that, the population mean of the health status proxy is stable, at approximately 1.6.

As we are matching the shape of health expenditure distributions, which often have a small number of extreme outliers, we trim the distribution. The nonparametric approach benefits from trimming the tails of the distribution and the data points that are dropped are not needed for a consistent estimator. We drop enrollees with zero yearly expenditures and whose expenditures fell in the top 20% of the entire sample of expenditures over all three plans.[17] Table 4 displays the resulting estimation sample. Over 10,000 employer-year observations remain in the risk parameter estimation sample.[18]

[16] It should be noted that we do not use this health status proxy measure in recovering the latent health status distribution. However, in analysis not reported, we do find that there is a high degree of correlation between this measure and our estimated latent health shock $\hat{\theta}$ for each individual.

[17] Additionally, a very small number of enrollees, less than 1% in each year, were dropped if their out-of-pocket expenditures were larger than their salaries.

[18] If there are nonnegligible number of zeros in m across the three years, two different approaches can be taken. The first approach, the one we have taken, is to assume a two-stage decision process. In the first stage, each potential

The GMM estimation procedure to recover utility parameters is implemented as follows. We explicitly describe the estimation procedure for the first moment, the mean. The other moments are constructed similarly.[19] The latent health status of each individual enrollee i, $\hat{\theta}_i$, is as specified in equation (5.2). Once we substitute in the data on observed health expenditure, income, insurance plan premiums, and reimbursement probabilities, that is, $w_{ijt} = (m_{ijt}, y_{it}, p_{jt}, \hat{f}_{jt})$, each individual i has an expression for her latent health status solely in terms of the utility parameters, γ. Equation (5.4) displays the mean of the health status distribution for the years 2002 and 2003, that is, $\mu_{\hat{\theta}_{02}}$, where N_{02} is the number of enrollees in 2002, as well as the corresponding definition of mean in 2003, $\mu_{\hat{\theta}_{03}}$.

$$\mu_{\hat{\theta}_{02}}(\gamma) = \sum_{\mathrm{i}=1}^{\mathrm{N}_{02}} \frac{1}{\mathrm{N}_{02}} \hat{\theta}_{\mathrm{i}}(\gamma), \qquad \mu_{\hat{\theta}_{03}}(\gamma) = \sum_{\mathrm{i}=1}^{\mathrm{N}_{03}} \frac{1}{\mathrm{N}_{03}} \hat{\theta}_{\mathrm{i}}(\gamma) \tag{5.4}$$

Let $h_1(w, \gamma_0)$ denote the first moment condition for our sample. Statistical inference using the GMM estimator is based on the property that $E[h_1(w, \gamma_0)] = 0$. For each of the four distribution moments, we have three moment conditions in our sample. These are the three pairs generated by comparing three years of data. The three sample moment conditions associated with the distribution mean are:

$$\begin{aligned}
\mathrm{h}_1(\mathrm{w}, \gamma) &= \left(\mu_{\hat{\theta}_{02}}(\gamma) - \mu_{\hat{\theta}_{03}}(\gamma)\right)^2 \\
\mathrm{h}_2(\mathrm{w}, \gamma) &= \left(\mu_{\hat{\theta}_{03}}(\gamma) - \mu_{\hat{\theta}_{04}}(\gamma)\right)^2 \\
\mathrm{h}_3(\mathrm{w}, \gamma) &= \left(\mu_{\hat{\theta}_{02}}(\gamma) - \mu_{\hat{\theta}_{04}}(\gamma)\right)^2.
\end{aligned}$$

Using the other THREE distribution moments, we generate NINE additional sample moment conditions for a total of 12 sample moment conditions. The GMM estimator minimizes the sum of all 12 sample moment conditions. This minimum value is found through a grid search over the possible values of the two utility parameters. Grids of varying size were used for $\hat{\gamma}_1 \in [1, 6]$ and $\hat{\gamma}_2 \in [1, 6]$. We substituted each combination of grid values for γ_1 and γ_2 into the sample moments and found the resulting sum of squared differences for that combination. The optimal $\hat{\gamma}$ is the grid combination of γ_1 and γ_2 with a sum of squared differences closest to zero.

We estimate coefficients of relative risk aversion for aggregate consumption, γ_1, and health, γ_2. Table 5 displays values of $\hat{\gamma}$ using increasingly fine grids over the parameter space. The standard errors were computed using the method described in the Appendix.

The resulting risk coefficients are in the range of [1.88, 1.98] for γ_1 and in the range of [3.12, 3.27] for γ_2.[20] Higher values in the γ_1 range tend to be associated with correspondingly higher values of γ_2. This result implies that individuals are more risk averse with respect to health status than to the aggregate consumption commodity. A slightly higher health risk coefficient implies that consumers are more risk averse with respect to their health, as it often cannot be regained once lost.

---

patient is endowed with a binary outcome. Either the patient is sick or not. If the patient is not sick, then she does not seek care and $m = 0$. If a patient is sick, then she gets a realization of θ, seeks care, and spends m as a function of θ, the parameters, and the covariates in the model. Using this two-step approach, the implicit assumption that we employ is that the distributions of θ given the occurrence of sickness do not vary over time. In our data, we find the proportion of zero ms does not vary over the three years, and this assumption appears to be a plausible one. The proportion of zeros in our sample is 11.3% in 2002, 10.7% in 2003, and 10.5% in 2004, which cannot be statistically distinguished in pairwise tests for all year combinations.

The second approach depends on corner solutions and a threshold level of θ^*, below which a consumer does not seek care. See the Appendix for a discussion of a general identification strategy that might capture the moment distributions of this approach.

[19] In recovering the utility parameters, we match the distributions using the first four moments of mean, variance, kurtosis, and skewness within each year.

[20] The aggregate consumption estimates are within the range found in the literature on consumption (see Gourinchas and Parker, 2002; Prescott, 1986; Shea, 1995).

TABLE 5 Estimated Risk Coefficients $\hat{\gamma}_1$, $\hat{\gamma}_2$

Grid Size	$\hat{\gamma}_1$	$\hat{\gamma}_2$
20	1.98	3.27
	(0.76)	(1.20)
40	1.88	3.12
	(0.86)	(1.35)
50	1.93	3.23
	(0.86)	(1.35)

Standard errors in parentheses.
Estimates from 200 bootstrap iterations.

FIGURE 2

Health status Distribution, 2002–2004

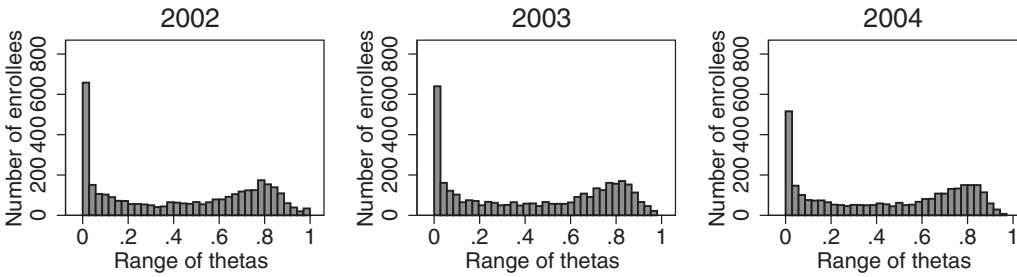

Recover the distribution of the latent θs. The final step of the estimation is to construct the latent $\hat{\theta}$ health status distribution. Substitute in $\hat{\gamma}_1$ and $\hat{\gamma}_2$ into equation (5.2) for an individual value of $\hat{\theta}_i$ for each individual enrollee in each year.

Figure 2 displays the estimated $g(\hat{\theta})$ distribution of health status for the years 2002–2004. Health status is on the horizontal axis, between 0 and 1. On the vertical axis is the number of consumers at each point in the interval. The distribution appears to be bimodal–many consumers in each year had very low values of θ, but there is a small clustering near the value of 0.8. This is a typical health distribution for a large general population–many healthy individuals but a small number of very sick individuals near the top of the distribution.

6. Results

□ **Adverse selection.** To test for sorting among plans, we break down the estimated latent health status distribution by plan. The HMO plan features generous coverage for network physicians and hospitals, but no coverage of out-of-network care. The PPO plans each have similar cost–sharing parameters, with nominal copayments for services in-network but lower coverage out-of-network after meeting a deductible. PPO1 has a network of providers of lower cost than PPO2. Figure 3 shows the distribution of latent health status is over all years in each of the three plans. The horizontal axis is the 0 to 1 range of the θ, and the vertical axis is the fraction of consumers in each plan corresponding to each θ value. A simple visual analysis shows that the HMO plan has a much larger fraction of very healthy individuals than the PPO plans–those with θ near zero. Over 20% of HMO consumers had a θ value less than 0.05, compared to the PPO plans with approximately 15% of consumers in this range. Both PPO1 and PPO2 plans show a larger clustering than the HMO plan above the health status value of 0.8.

FIGURE 3

Health Status Distribution, by plan

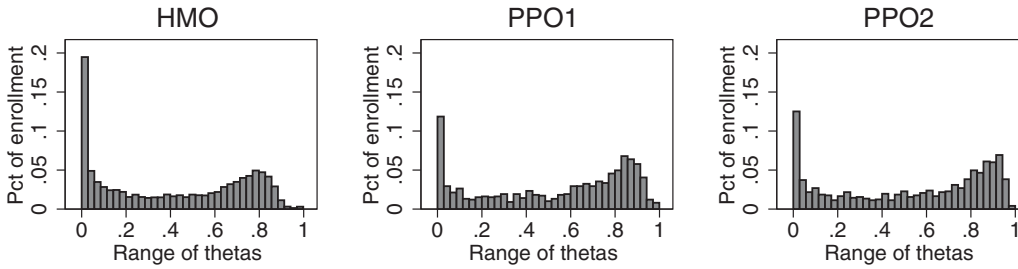

TABLE 6 Test For Sorting of θ Distribution Among Plans

Year	K-S Statistic	Outcome
H_0: HMO and PPO1 are from the **same continuous distribution**.		
H_A: HMO's **cdf is greater** than PPO1's cdf.		
2002	0.1948	Reject
2003	0.1778	Reject
2004	0.2347	Reject
H_0: HMO and PPO2 are from the **same continuous distribution**.		
H_A: HMO's **cdf is greater** than PPO2's cdf.		
2002	0.1638	Reject
2003	0.3036	Reject
2004	0.1864	Reject
H_0: PPO1 and PPO2 are from the **same continuous distribution**.		
H_A: PPO2's **cdf is greater** than PPO1's cdf.		
2002	0.0097	Do not reject
2003	0.1819	Reject
2004	0.0204	Do not reject

K-S = Kolmogorov–Smirnov.

To support the graphs' implications, we next examine which plans contain the higher tail (sicker) of the health status distribution using Kolmogorov–Smirnov (K-S) statistics.[21] The trade-off for consumers between an HMO plan and PPO plans is that PPO plans provide greater flexibility in provider choice in exchange for more cost sharing by consumers. Relatively healthy consumers who do not expect to insure more than a yearly checkup may be satisfied with the more limited network of providers in the HMO plan.

Table 6 tests for the direction of sorting within the plans. The first two sets of results test whether the HMO and PPO cdfs are from the same distribution, with the alternate hypothesis that the HMO cdf stochastically dominates the PPOs' cdfs. The K-S tests in Table 6 show the proportion of relatively healthy consumers is larger in the HMO plan than in either the PPO1 or PPO2 plan, by rejecting the null hypothesis. When comparing the PPO plans to each other in 2003, the hypothesis of equality of distributions between PPO1 and PPO2 is rejected in favor of PPO1 having a larger cdf. This means that distribution of consumers is healthier in the PPO1 plan. This hypothesis cannot be rejected for the other two years. This inability to reject may be expected given similar cost-sharing structures.[22]

[21] We also calculated K-S statistics that rejected the hypothesis that the HMO plans and the PPO plans have common distributions, without testing the direction of sorting.

[22] In an earlier working paper (Bajari, Hong, and Khwaja, 2006), we found no evidence of adverse selection using a different data set, the Health and Retirement Survey (HRS). However, it is our view that these findings were primarily driven by the data and not the methods. The HRS data places insurance plans in broad categories, for example, uninsured,

FIGURE 4

Estimated Overconsumption, 2002–2004

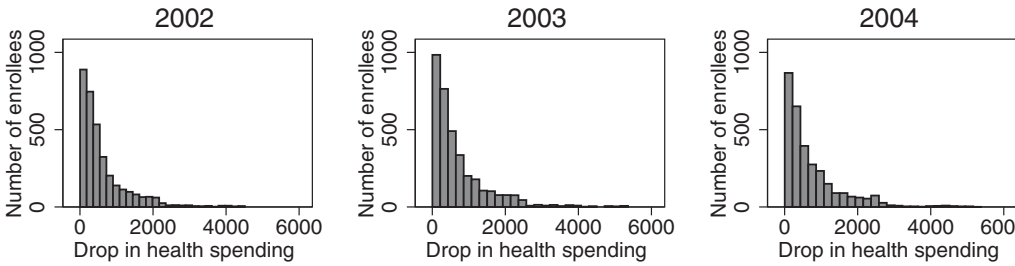

TABLE 7 Overconsumption as Percentage of Original Health Care Expenditure

Year	Mean	Median	Standard Deviation	*N*
2002	45.14	42.09	8.83	3433
2003	46.22	42.14	8.85	3533
2004	46.22	42.49	10.04	3100

□ **Moral hazard.** One factor cited in rising health care costs is "overconsumption," when consumers are insulated from costs in their health care choices. Our measure of moral hazard is based on a counterfactual which allocates to the consumer the same resources previously consumed in both health care and aggregate consumption, but now allows the consumer to choose a new allocation.

We find the magnitude of overconsumption in our data is substantial. Figure 4 shows the resulting distributions of moral hazard overconsumption estimates for 2002–2004. The horizontal axis is the drop in health spending from the original observed expenditure to the counterfactual expenditure. The vertical axis is a count of the number of consumers at each overconsumption level. Approximately half of consumers overconsume between $0 and $500 in each year, after which the distribution drops rapidly with a long right tail finishing above $3000 in overconsumption.

Table 7 lists summary statistics by year for overconsumption as a percentage of a consumer's original health care expenditure. The median overconsumption is over 40% of the original choice of health expenditure. This level of moral hazard is slightly smaller than that found in Keeler and Rolph (1988) based on the RAND Health Insurance Experiment (HIE), the gold standard in the health economics literature. Keeler and Rolph (1988) used the HIE to show that going from no insurance to full coverage results in a 50% increase in health expenditures. However, although the HIE did give lump-sum payments to cover worst-case scenarios from participating in the experiment, our counterfactual's income neutrality may explain why consumers in our sample reduce expenditures by slightly less than those in the HIE. Our results suggest that failing to control for income effects may overstate the level of moral hazard.

employer-provided, Medicare, VA-Champus, etc. In particular, the different employer-provided insurance plans were combined into one category, whereas employers typically offer a choice of multiple plans. The overwhelming majority of observations were for employer-sponsored insurance, 72.5%, the uninsured were 8.7%, the remaining categories each were less than 20%. That data also lacked detailed measures of plan characteristics. Thus, it was difficult to tease out adverse selection in the employer plans between less generous, for example, HMO plans, and more generous, for example, PPO plans. The current data focuses exclusively on one employer and has much better measures of variation between employer-provided plans, which helps us surmount the shortcomings of the previous data set. Somewhat presciently in that version (p. 30) we had stated, "It is possible that we are unable to find evidence of adverse selection because our insurance categories are very broad (e.g., in our data, employer-provided insurance is one category whereas employers typically offer a choice of multiple plans to employees). Conceivably, evidence for adverse selection may be found if an examination was done at a more detailed level, for example, across different kinds of employer-provided plans."

7. Conclusion

■ We present a new approach to measuring moral hazard and adverse selection, based on a model of health care demand in the presence of unobserved heterogeneity in health status. This approach does not also require information on plan choice, in contrast to existing approaches. We make three key contributions to the existing literature. First, we disentangle both moral hazard and adverse selection through estimation of a latent health status distribution, using detailed claims-level data. We compute a measure of moral hazard by performing a counterfactual to isolate the effect of insurance on health expenditure. Differences in latent health status distributions across plans measure adverse selection. The second contribution is our semiparametric estimation method. We estimate insurance reimbursement schedules nonparametrically, which has two advantages: (i) it allows for complex insurance plans including copays, deductibles, and other nonlinear features, and (ii) nonparametric estimation of a conditional reimbursement probability relaxes a common assumption in previous literature by incorporating the reality that plans are complex and consumers have difficulty predicting reimbursement in every state of the world. Finally, the third contribution is that we use a single, relatively weak identification assumption–that the distribution of health shocks is invariant over a short, consecutive time span. This assumption is reasonable in many commonly available large claims-level data sets.

Although our proposed semiparametric method provides a more flexible and robust alternative for analyzing adverse selection and moral hazard, there are caveats. We assume that the utility function is separable in the aggregate consumption commodity and health care. Although this captures risk aversion in health status, it rules out more flexible interactions between aggregate consumption and health status (see e.g., Viscusi and Evans, 1990). However, Spence and Zeckhauser (1971) and Blomqvist (1997) use a similar specification, and Campo et al. (2011) also require similar restrictions on utility in an auctions context. In spite of these limitations, our research is novel in that it develops a tractable estimation procedure under parsimonious parametric assumptions to simultaneously examine adverse selection and moral hazard. Our research is also important as it provides a framework for similar analysis in other contexts where distortions exist due to asymmetric information, especially when the data is cross sectional.

References

Akerlof, G. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism,"*Quarterly Journal of Economics*, Vol. 84(3) (1970), pp. 488–500.

Arrow, K.J. "Uncertainty and the Welfare Economics of Medical Care."*American Economic Review*, Vol. 53 (1963), pp. 941–973.

Bajari, P., Hong, H., and Khwaja, A. "Moral Hazard, Adverse Selection and Health Expenditures: A Semiparametric Analysis." NBER Working Paper no. 12455, 2006.

Bajari, P., Hong, H., Park, M., and Town, R. "Estimating Price Sensitivity of Economic Agents Using Discontinuity in Nonlinear Contracts." NBER Working Paper No. 17643, 2013.

Blomqvist, A.G. "Optimal Nonlinear Health Insurance."*Journal of Health Economics*, Vol. 16(3) (1997), pp. 303–321.

Bowman, A.W. and Azzalini, A. *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*. New York, NY: Oxford University Press, 1997.

Bundorf, M.K., Levin, J., and Mahoney, N. "Pricing and Welfare in Health Plan Choice,"*American Economic Review*, Vol. 102(7) (2012), pp. 3214–3248.

Campo, S.M., Guerre, E., Perrigne, I., and Vuong, Q. "Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders," *Review of Economic Studies*, Vol. 78(1) (2011), pp. 112–147.

Cardon, J.H. and Hendel, I. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey,"*RAND Journal of Economics*, Vol. 32(3) (2001), pp. 408–427.

Carlin, C. and Town, R."Adverse Selection, Welfare and Optimal Pricing of Employer Sponsored Health Plans." Working Paper, 2010.

Chiappori, P.-A. and Salanié, B. "Testing Contract Theory: A Survey of Some Recent Work." In L. Hansen, M. Dewatripont, and P. Turnovsky, eds., *Advances in Economics and Econometrics—Theory and Applications, Eighth World Congress*. Econometric Society Monographs, New York, NY: Cambridge University Press, 2003.

Courty, P. and Li, H. "Sequential Screening," *Review of Economic Studies*, Vol. 67 (2000), pp. 697–717.

CUTLER, D.M. and ZECKHAUSER, R.J. "The Anatomy of Health Insurance." In A.J. Culyer and J.P. Newhouse, eds., *Handbook of Health Economics*, Vol. 1A Amsterdam, New York: Elsevier, 2000.

DAI, C., LEWIS, T., and LOPOMO, G. "Delegating Management to Experts."*RAND Journal of Economics*, Vol. 27(3) (2006), pp. 403–420.

DALTON, C.M. "Estimating Demand Elasticities Using Nonlinear Pricing." *International Journal of Industrial Organization*, (2014), Forthcoming.

EINAV, L., JENKINS, M., and LEVIN, J. "Contract Pricing in Consumer Credit Markets."*Econometrica*, Vol. 80(4) (2012), pp. 1387–1432.

EINAV, L., FINKELSTEIN, A., RYAN, S.P., SCHRIMPF, P., and CULLEN, M.R. "Selection on Moral Hazard in Health Insurance."*American Economic Review*, Vol. 103(1) (2013), pp. 178–219.

GOURINCHAS, P.O. and PARKER, J. "Consumption Over the Life Cycle."*Econometrica*, Vol. 70(1) (2002), pp. 47–89.

GUERRE, E., PERRIGNE, I., and VUONG, Q. "Optimal Nonparametric Estimation of First-Price Auctions."*Econometrica*, Vol. 68(3) (2000), pp. 525–574.

HANDEL, B. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts."*American Economic Review*, Vol. 103(7) (2013), pp. 2643–2682.

KEELER, E. and ROLPH, J. "The Demand for Episodes of Treatment in the Health Insurance Experiment."*Journal of Health Economics*, Vol. 7 (1988), pp. 337–367.

KEELER, E., ROLPH, J., NEWHOUSE, J., and PHELPS, C. "Deductibles and the Demand for Medical Care Services: The Theory of a Consumer Facing a Variable Price Schedule under Uncertainty." *Econometrica*, Vol. 45 (1977), pp. 641–655.

KHWAJA, A. "Estimating Willingness to Pay for Medicare Using a Dynamic Life-Cycle Model of Demand for Health Insurance."*Journal of Econometrics*, Vol. 156 (2010), pp. 130–147.

PAULY, M. "The Economics of Moral Hazard: Comment." *American Economic Review*, Vol. 58 (1968), pp. 531–536.

PRESCOTT, E. "Theory Ahead of Business Cycle Measurement." *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 10(4) (1986), pp. 9–22.

ROTHSCHILD, M. and STIGLITZ, J.E. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, Vol. 90(4) (1976), pp. 630–649.

SHEA, J. "Union Contracts and the Life-Cycle/Permanent-Income Hypothesis." *American Economic Review*, Vol. 85(1) (1995), pp. 186–200.

SPENCE, M. and ZECKHAUSER, R. "Insurance, Information, and Individual Action." *American Economic Review*, Vol. 61(2) (1971), pp. 380–387.

VISCUSI, K. and EVANS, W. "Utility Functions that Depend on Health Status: Estimates and Economic Implications."*American Economic Review*, Vol. 80 (1990), pp. 353–374.