

Moral Hazard, Adverse Selection and Health Expenditures: A Semiparametric Analysis*

Patrick Bajari, Department of Economics, University of Minnesota and NBER
Han Hong, Department of Economics, Duke University
Ahmed Khwaja, Fuqua School of Business, Duke University

JEL Classification: C14, D82, I11

Keywords: Adverse Selection, moral hazard, health insurance, semiparametric estimation

Abstract

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. We empirically study this question by using data from the Health and Retirement Study to estimate a structural model of the demand for health insurance and medical care. Using a two-step semi-parametric estimation strategy we find significant evidence of moral hazard, but not of adverse selection.

* We have benefited from the comments of Amy Finkelstein and participants of the Conference on Structural Models in Labor, Aging and Health, and the Seventeenth Annual Health Economics Conference. 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.

I. Introduction

There is a large theoretical literature that predicts that asymmetric information in insurance markets may generate inefficient outcomes due to adverse selection and moral hazard (see e.g., Arrow 1963, Pauly 1968, Akerlof 1970, Zeckhauser 1970, Spence and Zeckhauser 1971, Spence 1973, Pauly 1974, Rothschild and Stiglitz 1976, Wilson 1977, 1980). Rothschild and Stiglitz (1976) develop a model in which adverse selection is present because individuals have private information about their health status. They demonstrate that insurers may inefficiently ration health care by capping benefits for the healthiest consumers. Similarly Pauly (1968) shows that moral hazard may be present in health insurance markets because consumers do not bear the full cost of health care expenditures.¹

The predictions of theoretical models of insurance markets can depend quite delicately on whether adverse selection or moral hazard is more important. Because of the complexity of insurance markets, theoretical models frequently emphasize one of these distortions at the expense of the other. Furthermore, optimal policy depends crucially on which of these distortions is most important. Thus, empirically assessing whether moral hazard or adverse selection is important is useful for guiding both theory and public policy.

It is well recognized that it is empirically difficult to distinguish between moral hazard and adverse selection and consequently there is little consensus on which of these

¹ The empirical evidence on the existence of asymmetric information in insurance markets is mixed. For example, recent work by Cawley and Philipson (1999) for life insurance markets; Cardon and Hendel (2001) for health insurance markets; and Chiappori and Salanie (2000) for automobile insurance markets, finds no evidence of asymmetric information. On the other hand, Finkelstein and Poterba (2004) find such evidence in annuity markets. 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 generosity of the insurance contract and adverse outcomes, and moral hazard based on the coinsurance elasticity of the demand for medical care.

two sources of inefficiency is more important. A common method to detect asymmetric information is to examine the correlation between risk outcomes and a measure of the generosity of a contract. However, as pointed out by Chiappori and Salanié (2003), among others, under moral hazard the generosity of the contract will lead to adverse risk outcomes while under adverse selection the causality is reversed, leading to observational equivalence between the two hypotheses.² Abbring, Chiappori, Heckman and Poinquet (2003) have suggested that one could exploit the dynamic consequences of experience rating in insurance markets to distinguish between adverse selection and moral hazard. However, U.S. health insurance markets are regulated to restrict experience rating, which precludes this proposed empirical strategy.³

In this paper, we propose an alternative approach to the assessing the importance of moral hazard and adverse selection in health insurance markets. Our approach is based on estimating a structural model of consumer demand for health insurance and medical utilization. The structural model is consistent with the theoretical models proposed by Spence and Zeckhauser (1971) and Blomqvist (1997) and allows for both adverse selection and moral hazard. In the model, risk averse agents have preferences over a composite commodity and health status, and maximize utility subject to a budget

² A test of asymmetric information based on the positive correlation between the generosity of the contract and adverse risk outcomes may also breakdown if there is heterogeneity in risk preferences (De Meza and Webb 2001). For example, if individuals who are highly risk averse also put more effort in to lowering their risk, and vice versa those who are less risk averse take fewer precautions, then there could be a negative correlation between the generosity of the contract and risk outcomes. Evidence of heterogeneity in risk preferences has been found by Finkelstein and McGarry (forthcoming) in the long term care insurance market and Cohen and Einav (2005) in the automobile insurance market. Chiappori, Jullien, Salanié and Salanié (forthcoming) develop a non parametric test to detect the presence of asymmetric information based on the correlation between the generosity of the contract and risk outcomes that addresses this limitation. However their test does not distinguish between adverse selection and moral hazard. Finkelstein and Poterba (2006) develop a test for adverse selection that avoids the limitation of heterogeneity in risk preferences. Their test is based on using data on observable characteristics of individuals that are correlated with the outcomes but are not used by insurers in pricing contracts.

³ Another alternative would be to use field experiments to investigate the presence of asymmetric information. Karlan and Zinman (2005) provide an example of such an experiment in credit markets in South Africa. However the feasibility of implementing such experiments in the context of health insurance, especially in the U.S., is questionable.

constraint. Agents have private information about their (latent) health status which is unobserved to the insurer leading to adverse selection. The agents do not pay the full costs of their health care coverage and therefore face a moral hazard problem. We estimate our model using a combination of publicly available and confidential data from the Health and Retirement Study (HRS), which is a nationwide sample that contains information on health insurance plans, respondent location, medical expenditures and reimbursements from insurance plans.

The model is estimated semiparametrically. Our econometric strategy imports techniques from the empirical auctions literature to the study of insurance markets. Following Campo, Guerre, Perrigne and Vuong (2003), utility is specified using a standard functional form. This is convenient for comparing our estimates of risk aversion to existing estimates in the literature. However, we estimate the distribution of private information nonparametrically. Guerre, Perrigne and Vuong (2000) argue that this is important because theory provides little guidance about which parametric distributions for latent health shocks are a priori most plausible.

To foreshadow our results, first, we find that our model generates many results that are consistent with the existing empirical literature. Our estimates of risk aversion are consistent with much of the previous literature. Also, we find that the median elasticity of the consumption of health with respect to the co-payment rate is consistent with the well known RAND Health Insurance study (Manning et al 1987, Newhouse 1993) of the late 1970's.

Second, we find that the elasticity of health care usage with respect to the co-payment rates is highly nonlinear. This usage elasticity is several times higher for low levels of health consumption than for high levels of health consumption. We find that the

“sickest” consumers respond very inelastically to changes in co-payment rates compared to “healthy” consumers.

Third, we propose a nonparametric test for adverse selection. Theoretical models predicting separating equilibrium suggest that consumers should sort across different insurance categories based on their health status, which is private information to the consumer. Thus, we test whether our nonparametrically estimated distributions of health shocks differ across insurance categories that vary in their premium and coinsurance rate. With the exception of consumers who self insure, we find no statistically significant difference in health status across major insurance categories. This indicates a lack of evidence supporting models of the insurance market that predict a separating equilibrium, and which are commonly used in the theoretical literature.

There has been a large empirical literature on the effects of health insurance on medical utilization (see Zweifel and Manning 2000 for an excellent review). The gold standard of this work is considered to be the RAND HIE of the late 1970’s (Manning et al 1987, Newhouse 1993). We note that our results agree with this earlier work for the median consumer. However, there is considerable variation in these elasticities between the most and least healthy consumers. We argue that a nonparametric perspective is important because a linear econometric model masks substantial variation in the distribution of elasticities.

There is a growing body of empirical work on the structural estimation of models of medical utilization and health insurance choice (e.g., see the important work of Cameron et al. 1988, Gilleskie 1998, Harris and Keane 1999, Cardon and Hendel 2001, Vera-Hernandez 2003, Blau and Gilleskie 2003, Khwaja 2001, 2005). However, we differ from this earlier literature in three ways. First, this earlier work uses methods from

discrete choice estimation and relies on parsimoniously specified parametric models. However, akin to the empirical literature on auctions, our estimation strategy is semiparametric. Second, we allow for both adverse selection and moral hazard due to asymmetric information in the estimation procedure. Third, the earlier literature, though it controls for the endogeneity of selection in to insurance plans, assumes that prices are exogenous. This contradicts the theoretical literature which predicts consumers sort between plans based on their private information about health status. We propose an instrumental variables strategy that accommodates not only selection in to insurance plans but also the price endogeneity of these contracts.

The rest of the paper is organized as follows. The model is presented in section II. Section III discusses the data, with the estimation and identification strategy discussed in section IV. Results follow in section V, with examination of adverse selection and moral hazard in section VI. Section VII concludes.

II. Model

II.A. Overview

We specify a model of endogenous consumer demand for health insurance and medical utilization that allows for heterogeneity in the distribution of latent health status. Individuals have private information about their latent health status at the time of making choices about health insurance and consumption of medical care. There is asymmetric information in that insurers know the distribution of the latent health status but do not have the same information that the individuals have. Insurers however recognize the existence of this asymmetric information and formulate their contracts appropriately. In our empirical analysis we allow for the endogeneity of the insurance contract with regard to the distribution of latent health status. We focus on estimating the demand side

parameters while flexibly accommodating the endogeneity of the supply of health contracts. Given the complex nature of the health insurance market, it is not often clear what the economic objective of insurers is, i.e., whether insurers maximize profits (e.g., for-profit insurance plans) or pool social risk at the expense of profits (e.g., Medicare) or provide fringe benefits in kind (e.g., employer provided insurance). Hence we do not explicitly model the insurer's problem. However, our model is robust to adverse selection and moral hazard that may be generated due to asymmetric information that individuals have about their latent health status relative to insurers.

II.B. Model Structure

II.B.1. The Timing Convention

Individuals in the model are assumed to make decisions about the purchase of health insurance and medical care in a staggered fashion. Following Cardon and Hendel (2001), and Khwaja (2001, 2005) there are two time periods in the model. Given a menu of insurance options, individuals in the first time period make a choice about an insurance plan. In the second time period, conditional on the first period insurance choice and a realization from the distribution of latent health status, the individuals make a choice about medical utilization.

II.B.2. The First Time Period: Insurance Choice

The individuals face a menu of choices from a set of available insurance plans (D). Each individual has private information about their latent health status. Given their age, each individual self selects in to one of the insurance plans $d \in D$ based on his or her private information about latent health status. Our estimation strategy (described below) allows for adverse selection based on the distribution of latent health status even

though we do not explicitly model the process through which individuals self-select in to insurance plans.

II.B.3. The Second Time Period: Medical Utilization Choice

Conditional on the insurance choice in the first time period and a realization from the distribution of latent health status, each individual makes a decision about medical utilization in the second period. The individuals choose the level of medical utilization that maximizes their utility from health status and the consumption of a composite good. The utility function and the budget constraint are described below.

II.B.4. The Utility Function

Following Spence and Zeckhauser (1971) and Bloomqvist (1997) we specify the consumer's utility function, $U(c, m-\theta; \gamma)$, to depend on the level of composite good consumption c , the amount of medical utilization m , the consumer's latent health status θ , and the parameters γ that characterize the utility function of the consumer. As in Cardon and Hendel (2001), medical utilization is assumed to be a perfect substitute for health shocks in the individual's utility function. This is a restrictive assumption but as found previously by Cardon and Hendel it captures the essential features of the data well. Also this allows for the preventive aspects of medical care, i.e., people incur medical expenditures even in good health for preventive purposes. Both m and θ are assumed to be expressed in terms of monetary units. Therefore only the difference between m and θ directly enters the utility function. In particular we specify the utility function to be separable in c and $(m-\theta)$, and take the form,

$$U(c, m - \theta; \gamma) = \frac{1}{1 - \gamma_1} c^{1 - \gamma_1} + \frac{\gamma_2}{1 - \gamma_3} (m - \theta)^{1 - \gamma_3} . \quad (1)$$

The utility function allows for risk aversion in wealth (see e.g. Hubbard, Skinner and Zeldes 1995, Gertler and Gruber 2002) through the parameters on the composite

consumption good as well as in health status (see e.g., Khwaja 2001, 2005) through the parameters on the difference between medical utilization and latent health status.⁴ The budget constraint faced by the consumer is

$$c = y - p - z(m),$$

where y is an exogenously given level of income, p denotes the premium or the fixed cost of participating in the health insurance policy, and $z(m)$ denotes the out of pocket expenditure or the co-payment of the individual. Alternatively, the reimbursement scheme used by the insurance plan is that $\$(m - z)$ will be reimbursed to the consumer if the consumer incurs $\$m$ in medical expenses. We assume that the insurer specifies the reimbursement schedule $(m - z(m))$ prior to the realization of an individual's health shock θ . The fact that the reimbursement schedule is only a function of the medical utilization m and does not directly depend on the health status realization θ creates a moral hazard problem. After the realization of the health status θ , the consumer chooses the level of health services m as a function of θ , $m(\theta)$, to maximize his or her utility. We further assume that m is a non-decreasing function of θ . Whenever $m(\theta)$ is strictly positive the following first order condition holds,

$$U_m(c, m - \theta; \gamma) - U_c(c, m - \theta; \gamma)[z'(m)] = 0. \quad (2)$$

Equation (1) provides the condition for the optimal choice of medical utilization that maximizes an individual's utility conditional on the insurance status and the health shock realization. This is the standard marginal rate of substitution (MRS) rule for allocating income between the composite commodity and medical care. It states that the ratio of marginal utilities from medical care and the composite commodity should be equated to

⁴ We abstract away from heterogeneity in risk preferences and focus on testing for asymmetric information through direct measures of adverse selection and moral hazard elicited by our structural model. In principle our method can be extended to include preference heterogeneity. A reason we did not implement this was because of the nature of our data, which in particular limited the number of instruments and hence the number of endogenous parameters we could estimate.

their respective prices. Intuitively, conditional on the insurance contract an individual equates the marginal benefit of medical care measured in terms of improvement in utility to the marginal cost of medical utilization in terms of out of pocket expenditures using the composite commodity as the numeraire good. This optimality condition relates the unobserved health status θ to the observable medical utilization of the consumer m . Under the previous specification, this optimality condition becomes

$$\gamma_2(m - \theta)^{-\gamma_3} = c^{-\gamma_1}(z'(m)). \quad (3)$$

We use this condition as a basis for developing our identification and estimation strategy which we describe in greater detail in section IV.⁵

III. Data

We estimate our model using data from the Health and Retirement Study (HRS). The HRS is a nationally representative sample of men and women born between 1931 and 1941 and their spouses or partners, who could be of any age (see Juster and Suzman 1995 for an excellent and detailed review of the data, and also <http://hrsonline.isr.umich.edu>). At the initiation of the survey, in 1992 (wave 1), 12,652 people from 7,607 households were asked questions about economic and demographic characteristics and various life cycle choices. The study also included questions about respondents' health insurance choices and their medical expenditures. The HRS oversampled blacks, Hispanics, and residents of Florida. We use data on individuals from wave 3 (1996) of the HRS because this is the most recent wave for which we have data on out of pocket and total medical expenditures.⁶ The publicly available data is supplemented by confidential data on location of residence to help create instrumental

⁵ Such an approach to estimating structural parameters using first order conditions also has a long tradition in labor economics, e.g., Heckman and MaCurdy (1980). Recent work by Sieg (2000) also uses an approach based on first order conditions that involve derivatives of unknown functions.

⁶ We are grateful to Dan Hill at HRS for providing us with this data.

variables used in our analysis. The variables used in our analysis are described in detail in appendix I and summary statistics are in table 1.

In the empirical analysis we specify the set of insurance choices in the following way. Given the nature of our data, and the institutional features of the U.S. health insurance market (primarily that almost all individuals 65 or older have access to Medicare insurance) the insurance choice set (D) depends on whether individuals are younger than 65, or 65 or older. For individuals younger than 65 D consists of the following choices: employer provided health insurance, Veterans Administration/Champus insurance, insurance through own business for self-employed, privately purchased insurance, and no insurance (uninsured).⁷ For individuals 65 or older D consists of: employer provided health insurance, Veterans Administration/Champus insurance, insurance through own business for self-employed, Medicare insurance with or without private Medigap insurance.⁸ In the rest of the paper we use the term insurance plan and insurance category interchangeably. It is well known that different kinds of insurance plans may exist within each of the insurance categories included in the choice set but our data does not permit a more disaggregated analysis.

In wave 3 (1996) of the HRS we have data on 10,030 individuals in 6,351 households. Since the measure of income (in particular non wage income) is only available at the household level we compute the other variables in the budget constraint at the household level as well. We create household level variables by combining

⁷ It should be noted that the choice of health insurance is closely related to the employment decision of the individual. Our model does not include an employment decision in the interest of simplicity and due to the computational burden this would place on the estimator that we develop. An extensive review of the relationship between availability of health insurance coverage and the labor market decisions of individuals is provided by Currie and Madrian (1999) and Gruber (2000).

⁸In our sample 316 or 8.66% of the individuals are uninsured, 2644 or 72.48% have employer provided insurance, 81 or 2.22% are self-employed with insurance through their own business, 123 or 3.37% are on VA/Champus, 255 or 6.99% have privately purchased insurance, and 229 or 6.28% are on Medicare.

information on adults in the same household. Since the budget constraint is calculated for households, the model is estimated at the household level. Therefore in the rest of the paper, by individuals we mean individual households.

The data are trimmed for our analysis in the following way: (1) Observations where household income, household insurance premium, or household out of pocket medical expenditure was missing were dropped. This left a sample 4645 observations. (2) Observations where household out of pocket medical expenditure exceeded household total medical expenditure were dropped. This reduced sample size to 4540, a drop of approximately 2%. (3) Observations where household income was less than the sum of household insurance premium, and household out of pocket medical expenditure were dropped. This reduced sample size to 4412, a drop of approximately 3% of remaining observations. (4) Observations where the insurance category was Medicaid were dropped. This reduced the sample size to 4155, a drop of approximately 6% of remaining observations. (5) Outliers for household total medical expenditure were dropped. Outliers were observations with values of household total medical expenditure below the second percentile or above the ninety seventh percentile. This reduced the sample size to 3935, a drop of approximately 5% of remaining observations. (6) Observations where an individual was both on Medicare and younger than 65 were dropped. This reduced the sample size to 3724, a drop of approximately 5% of remaining observations. (7) Observations where the estimated derivative of out of pocket medical expenditure with respect to household total medical expenditure, was negative were dropped (for reasons that are further explained when describing the estimation strategy in section IV). This reduced the sample size to 3648, a drop of approximately 2% of remaining observations.

IV. Estimation

IV.A. Two Step Estimation Strategy and Identification

We propose a semiparametric estimator to recover the parameters of the utility function, γ . The major advantage of this strategy is that we do not have to rely on parametric assumptions about the latent health distribution in estimating the parameters of the model. The key insight of our identification strategy, is that both the co-payment rate $z'(m)$ (and in turn the reimbursement schedule $(m - z(m))$), and the distribution of the health status θ can be non-parametrically identified using the optimality condition (equation 3) about the level of medical utilization. These non-parametric estimates can in turn be used to estimate the risk aversion parameters, γ_1 and γ_3 . Therefore, our identification strategy depends only on the specification of the utility function and on the validity of the economic hypothesis of utility maximization but not on the statistical hypotheses regarding the reimbursement schedule and the distribution of latent health status.

The estimation method proceeds in two steps. In the first step, we nonparametrically estimate the health insurance co-payment schedules using data on out of pocket medical expenditures and insurance choices. The identification assumption is that insurance plans may design reimbursement policies given their expectation about the distribution of latent health status (in anticipation of adverse selection) but once individuals opt for a particular plan the reimbursement schedule cannot further discriminate against particular individuals. In other words the reimbursement schedule cannot be made individual specific though it may be group specific. We use a local linear estimator to nonparametrically recover from the data the co-payment schedule conditional on insurance choice.

Given data on $i = 1, \dots, n$ individuals about their choice of insurance plans d_i (specified here as a 5 by 1 vector of dummies for 5 of the six insurance categories with privately purchased insurance as the omitted category), whether they live in an urban area w_i , their level of out of pocket medical expenditures z_i and the level of total medical expenditures m_i , we estimate the mean regression equation for $z(q)$, where $q = [m \ d \ w]$, using a local linear estimator. We use the specification, $z = E[z | Q = q] + \varepsilon = z(q) + \varepsilon$.

Letting $z(q) = \alpha(q)$ and $\frac{dz}{dq_j}(q) = \beta_j(q)$, (where j is an index for the elements of the

vector q) we estimate $(\hat{\alpha}(q), \hat{\beta}(q))$ as,

$$\arg \min_{\alpha, \beta} \sum_{i=1}^n \left(z_i - \alpha - \sum_{j=1}^7 \beta_j (Q_{ij} - q_j) \right)^2 \cdot K_h(Q_i - q). \quad (4)$$

In particular, h is a sequence of bandwidth parameters, α is the conditional expectation, and β is a vector of derivatives in the local linear estimator (see e.g., Fan and Gijbels 1996, pp. 298-299 for details). In the local linear regression we smooth over observations using Gaussian kernel weights (K) for insurance plan dummies and the urban/rural dummy. This is done to utilize information from other plans in estimating the contract as the number of observations is small for some of the plans (e.g., self-employed insurance category). Letting σ_j be the standard deviation of the j^{th} variable in q , and N be the standard Gaussian density, we define the weighting function as

$$K_h(Q_i - q) = \prod_{j=1}^7 \frac{1}{h} N\left(\frac{Q_{ij} - q_j}{h\sigma_j}\right)$$

The bandwidth is fixed to be 0.95. Following Bajari and Kahn (2005) we chose the bandwidth based on visual inspection of the estimates because as pointed out by them, the bandwidth size suggested by asymptotic theory in Fan and Gijbels (1996) is unlikely

to be reliable in the presence of 7 regressors. This specification captures any non-linearity in the co-payment schedule e.g., deductibles, maximum annual out of pocket limit, maximum life-time reimbursement limit (see e.g., Keeler, Newhouse and Phelps 1977 for the complications that arise when such non-linearities are ignored).

In the second step, we use the instrument variables (denoted by x) to identify and estimate both the utility function parameters γ and the distribution of the latent health status θ . The key identifying assumption is the independence between the instruments (x) and the health status shocks (θ) (instruments are described in greater detail below). Even though the health status shocks (θ) are unobservable they can be uniquely recovered from the observable medical utilization by inverting the optimality condition given by equation 3. To illustrate, given any candidate value for the utility function parameters γ , for each individual $i = 1, \dots, n$, the unobservable health status θ_i can be recovered from the observed medical utilization m_i using the consumer optimality condition (equation 3), where $z(m_i)$ is replaced by the estimated function $\hat{z}(m_i)$ recovered using the local linear estimator.⁹ Hence θ_i can be written as

$$\hat{\theta}_i = \varphi(m_i, p_i, y_i, \hat{z}(m_i), \gamma). \quad (5)$$

In particular, given our utility specification (equation 1) and under the additional assumption that $z'(m) \geq 0$, equation 5 is uniquely defined as

$$\varphi(m_i, p_i, y_i, \hat{z}(m_i), \gamma) = m_i - \left[\frac{1}{\gamma_2} (y_i - p_i + \hat{z}(m_i) - m_i)^{-\gamma_1} (\hat{z}'(m_i)) \right]^{-\frac{1}{\gamma_3}}.$$

The economic intuition for the condition $z'(m) \geq 0$ is that as the medical utilization increases the co-payment should not decrease.

⁹ For notational convenience we suppress the dependence of $\hat{z}(m_i)$ on the individual's choice of insurance plan d_i , and whether he or she lives in an urban area w_i .

Using the instrument variables x_i for each individual $i = 1, \dots, n$ we formulate a method of moment estimator where the parameters of the utility function γ and the median of the unconditional distribution of health status (μ_θ) are jointly estimated by minimizing the following objective function,

$$(\hat{\gamma}, \hat{\mu}_\theta) = \arg \min \left\| \frac{1}{n} \sum_{i=1}^n g(x_i) \left\{ \mathbb{I} \left(\varphi(m_i, p_i, y_i, \hat{z}(m_i), \gamma) \leq \mu_\theta \right) - 0.5 \right\} \right\|_n, \quad (6)$$

where $\|\cdot\|_n$ is a typical quadratic norm i.e., $\|x\|_n = x'W_nx$, for a suitably chosen weighting matrix W that is used in generalized method of moments estimation, and $g(x)$ is a set of functions that generate different functional forms for a given vector of instrumental variables x . The estimation procedure iterates between using equations 5 and 6 to obtain consistent estimates of the utility parameters γ and the distribution of the latent health status θ .

In the HRS data we observe that some individuals have zero out of pocket medical expenditures, e.g., because they did not consume medical care.¹⁰ Hence we use a median based moment condition instead of the conventional mean based moment condition because the former is more robust to censoring at the upper and lower tails of the conditional distribution of the observables (see e.g., Powell 1984, Hong and Tamer 2003). Our two step method of moments estimation procedure also helps us avoid formulating an explicit selection equation for the choice of health insurance plans, as would be the case if we used a likelihood based approach.

The first stage estimates of the implicit prices will converge at a nonparametric rate. However, as demonstrated below, the estimates of the structural parameters will converge at the standard parameter rate, i.e. $n^{1/2}$. This semiparametric estimation strategy

¹⁰ In our estimation sample 3.14% or 125 observations have zero OOP medical expenditure. The 125 observations are broken down among insurance categories as: no insurance-24, employer provided-73, VA/Champus-19, privately purchased insurance-3, Medicare-6.

is attractive because it imposes minimal a priori restrictions on our first stage estimates but still allows for a parametric convergence rate.

Our identification strategy makes use of instruments that provide exogenous variation in the characteristics of the health insurance plans. The instruments we use come from geographic variation in (i) the state level housing price index, (ii) the county level malpractice insurance component of the Geographic Practice Cost Index (GPCI) developed by the Medicare Payment Advisory Commission to reimburse medical practices that treat Medicare beneficiaries (see appendix for more details), and (iii) number of establishments in a county.¹¹

We assume that these instruments reflect county level variation in the costs of providing insurance and are independent of the distribution of latent health status of the consumers. Alternatively put, the identification assumption is that our instruments provide variation in the price of providing medical care relative to other goods but are uncorrelated with the latent health distribution. More specifically, higher costs of malpractice liability insurance should increase expenditures of medical providers directly, and indirectly may also increase the costs of providing medical care if it induces the practice of defensive medicine. Both of these would be reflected in the price of insurance. A higher state level housing price index should increase the costs of providing medical care by affecting construction and rental costs for medical infrastructure, as well as may affect the compensation to physicians because of higher costs of living. This in turn would be reflected in the price of insurance. A larger number of establishments in a county could affect the price of insurance in various ways. It could increase the employee

¹¹ We also tried other instruments that could potentially play the role of cost shifters in the supply of insurance (e.g., the number of employees in the firm that employs an individual, local unemployment rate, average local income, HMO penetration) but found these to have weak explanatory power. Hence these were not used in estimating the model.

risk pool and hence the lower the costs of providing insurance. Additionally it could provide more bargaining power to coalitions of employers in negotiating insurance contracts with insurers leading to lower prices of insurance. Also in equilibrium it may attract more insurers to the area leading to changes in the costs and prices of contracts due to aggregation effects.

In addition, a feature of the HRS data is that (almost) all individuals who are older than 65 are eligible for Medicare coverage. Thus Medicare eligibility provides exogenous variation in the set of insurance choices which potentially helps in the identification of the distribution of latent health status.

IV.B. The Relationship with Linear Instrumental Variables Estimation

Our estimation strategy is similar to the use of instrumental variables in conventional two stage least square estimation of demand equations. The unobserved health status θ is the analogue of the unobserved error term in a structural linear demand equation. In that case, once the parameters of the linear demand equation are estimated, the conditional or unconditional distributions of the error term can also be recovered non-parametrically by inverting the linear equation. This is similar to our procedure for recovering the distribution of θ by inverting the implicit function defined by equation 3. Once we have the estimates of the co-payment schedule, $\hat{z}(m)$, from the first stage, we can rewrite equation 3 as,

$$\gamma_2(m - \theta)^{-\gamma_3} = c^{-\gamma_1}(\hat{z}'(m)). \quad (7)$$

Given that we observe data on consumption (c) and medical expenditures (m), we recover utility parameters γ and the distribution of the latent health status θ in the following way. The instruments (x) generate exogenous shifts in the relative price of medical care that are uncorrelated with the health shock (θ). Thus intuitively we can hold

θ “fixed” while the relative prices are shifted by the instruments, which leads to co-movement in c and m . Such co-movements allow us to identify the utility parameters γ . Once we have estimated the utility parameters we can use equation 3 (or 7) to back out the distribution of health status (θ). 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 in equation 3 to derive the functional form of the demand equation.

Previous research, such as Cardon and Hendel (2001), that uses sophisticated discrete choice models assumes that health shocks are uncorrelated with characteristics of the insurance plan, such as price or reimbursement rates. Nonetheless, standard models of insurance markets would predict that unobserved health status should be correlated with price or other contract characteristics if there is a separating equilibrium. There is no obvious solution based on an IV strategy to the problem of endogeneity in a discrete choice framework. An advantage of our approach is that we are able to instrument in a fairly straightforward manner.

A second attractive feature of our estimator is that we are able to derive a nonparametric distribution of agent’s private information. There are very few papers that attempt to structurally estimate empirical models of contracts, with the notable exceptions of Ferall and Scherer (1999), Paarsch and Scherer (2000), and Cardon and Hendel (2001). Because of the complexity of these models, these researchers are forced to make fairly strong parametric assumptions about the distribution of private information. In the empirical auctions literature (e.g., Guerre, Perrigne and Vuong (2000) and Athey and Haile (2002)) it is commonly argued that theory provides very little guidance about the

appropriate parametric distribution for agent's private information. Furthermore, it is commonly argued that conclusions about distortions from informational rents and other distortions will be biased if ad hoc parametric assumptions about the distribution of private information are imposed on the model. Therefore, the literature has emphasized the importance of estimating private information under the weakest possible parametric restrictions. To the best of our knowledge, this is the first paper to semiparametrically estimate a structural model with adverse selection and moral hazard.

IV.C. Computation of Standard Errors

In this section we discuss how we compute the correct standard errors for the two step semiparametric estimator for the parameters γ and μ , where γ are the parameters in the utility function and μ are the nuisance parameters that characterize the median of the distribution of the health status shock θ .

The estimator that we use in this paper is based on moment conditions of the form

$$m_n(\hat{\gamma}, \hat{\mu}, \hat{z}(\cdot)) = \frac{1}{n} \sum_{i=1}^n h(x_i, \phi(w_i, \hat{z}(m_i), \hat{z}'(m_i), \hat{\gamma}), \hat{\mu}),$$

where x_i denote the instruments and w_i denote the data m_i, p_i, y_i and d_i , for each individual $i = 1, \dots, n$. For example, equation (6) gives the particular form of the moment condition $h(x_i, \phi(w_i, \hat{z}(m_i), \hat{z}'(m_i), \hat{\gamma}), \hat{\mu})$ we use in the estimation procedure. The GMM estimator $\hat{\alpha} = (\hat{\gamma}, \hat{\mu})$ is calculated using the quadratic norm

$$\min_{\gamma, \mu} Q_n(\alpha) \equiv m_n(\hat{\gamma}, \hat{\mu}, \hat{z}(\cdot))' W_n m_n(\hat{\gamma}, \hat{\mu}, \hat{z}(\cdot)).$$

The asymptotic distribution of $\hat{\alpha}$, as is usual for a two step GMM estimator, is given by

$$\sqrt{n}(\hat{\alpha} - \alpha) \xrightarrow{d} N(0, (G'WG)^{-1}(G'W\Omega WG)(G'WG)^{-1}),$$

where $G = \frac{\partial}{\partial \alpha} E m_n(\alpha, z(\cdot))$, $W_n \xrightarrow{p} W$, and Ω is the limiting asymptotic variance of

$$\sqrt{nm_n}(\gamma, \mu, \hat{z}(\cdot)).$$

To derive Ω , we use the theory developed in Newey (1994), who showed that Ω does not depend on the particular nonparametric method that is used to estimate $z(\cdot)$ and $z'(\cdot)$.

To describe the asymptotic distribution using the framework of Newey (1994), define

$$\bar{h}(z(m_i), z'(m_i), m_i; d_i) = E[h(x_i, \varphi(w_i, z(m_i), z'(m_i), \gamma), \mu) | m_i, d_i]$$

and

$$d_1(m_i; d_i) = \frac{\partial}{\partial z(m_i)} \bar{h}(z(m_i), z'(m_i), m_i; d_i),$$

and

$$d_2(m_i; d_i) = \frac{\partial}{\partial z'(m_i)} \bar{h}(z(m_i), z'(m_i), m_i; d_i).$$

Then following Newey (1994), under suitable regularity conditions we can write, up to a term that converges to 0 in probability, $\sqrt{nm_n}(\gamma, \mu, \hat{z}(\cdot))$ as

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n ((\delta_1(m_i, d_i) + \delta_2(m_i, d_i))(z_i - z(m_i)) + h(x_i, \varphi(w_i, z(m_i), z'(m_i), \gamma), \mu)),$$

where we define

$$d_1(m_i; d_i) = \delta_1(m_i; d_i)$$

and

$$\delta_2(m_i; d_i) = -\frac{1}{f(m_i | d_i)} \frac{\partial}{\partial m} [d_2(m_i, d_i) f(m_i | d_i)].$$

Therefore

$$\Omega = Var((\delta_1(m_i, d_i) + \delta_2(m_i, d_i))(z_i - z(m_i)) + h(x_i, \varphi(w_i, z(m_i), z'(m_i), \gamma), \mu)).$$

Given that we know the form of the limiting variance, the next and final step is to obtain a consistent estimate of Ω , the asymptotic variance. For this we need to obtain consistent

estimates of the elements of the asymptotic linear influence function. Each of these components is related to the conditional expectation function,

$$\bar{h}(z(m_i), z'(m_i), m_i; d_i) = E[h(x_i, \varphi(w_i, z(m_i), z'(m_i), \gamma), \mu) | m_i, d_i],$$

which can be estimated nonparametrically and consistently. We obtain nonparametric estimates of the linear influence functions and then estimate the asymptotic standard errors by the empirical sum of the outer product of the estimated influence function.

V. Estimates of Model Parameters

We next discuss the results from the two step semiparametric estimation of the model presented in section III. Table 2 reports the summary statistics for the first stage local linear regression of the co-payment schedule $z(m)$. The omitted category is privately purchased insurance and for individuals in this baseline category the average out of pocket expenditure is \$1,676 (all values are in 1996 dollars). We find that uninsured individuals on average incur an out of pocket expenditure of \$1,333,¹² which is higher than for those with employer provided insurance who incur an average expenditure of \$1,133 but lower than the average expenditure of \$1,892 for self employed individuals with insurance through their own business. Individuals on Medicare incur an average expenditure of \$1,096, while those on VA/Champus have an average expenditure of \$914, which is the lowest among the insurance categories. At the mean a 1% increase in total medical expenditure leads to a 11% increase in out of pocket expenditures. Those who live in an urban area (defined as central counties of metropolitan areas of population of one million or more) incur an average lower expenditure of \$56. The differences in out of pocket expenditures across plans may be explained by differences in co-payment rates (i.e., incentives of the insurance plan) as well as differences in medical utilization, e.g.,

¹² This is the difference between the constant and the parameter on the dummy variable for insurance category 1.

the individuals who do not have insurance (and are not reimbursed) have lower average expenditures than those who are self-employed because they seek less medical care. We also find that the individuals in the self employed insurance category have the largest variation in their out of pocket expenditures, implying that they bear the greatest risk, whereas those in the employer provided category bear the least risk.

Figure 1 shows that the co-payment schedule has a concave shape. Figure 2 provides a description of the gradient of the co-payment schedule. The gradient is non-negative almost everywhere implying that the co-payment does not decrease as total medical expenses (or medical utilization) increase, i.e., $z'(m) \geq 0$. This is important for the implementation of our two stage estimation procedure.¹³ Figures 1 and 2 imply that the out of pocket expenditures are an increasing function of the total medical expenditures even though the co-payment rate decreases as total medical expenditures increase. These figures demonstrate that the expected co-payment rate is a nonlinear function of expenditures and therefore it is advantageous to model an agent's co-payment rates flexibly using a nonparametric framework. Figure 2 suggests that at the highest levels of expenditure, the co-payment rate is significantly smaller than at lower levels of expenditure. Our estimates suggest that this variation in co-payment rates occurs within a plan, rather than between plans as suggested by models predicting a separating equilibrium in insurance markets. Such a contract would be consistent with behavior if the incentives for individuals to economize on discretionary health expenditures at the margin are the least at the lowest levels of total medical expenditures.

¹³As described in the data section, in estimating the structural parameters we dropped 76 observations or 2% of the sample for which the derivative of the copayment schedule was negative.

Table 3 presents the second step estimates of the utility parameters. To our knowledge we are the first to estimate the coefficient of relative risk aversion with respect to health as distinct from that for aggregate consumption. Associated with this we also estimate the utility weight for consumption of health relative to aggregate consumption. The coefficient of relative risk aversion for the aggregate consumption commodity (γ_1) is estimated to be 0.85, while that for the consumption of health care (γ_3) is 1.52. Thus individuals are more risk averse with respect to health status than the aggregate consumption commodity. These numbers are broadly within the range of the estimates found in the literature on consumption (see e.g., Zeldes 1989, Shea 1995, Hansen and Singleton 1982, Gourinchas and Parker 2002). The utility weight (γ_2) on the consumption of health status relative to the aggregate consumption commodity is estimated to be 1.37. The individuals value consumption of health status more than they do the consumption of the aggregate consumption commodity. Hence an individual would have to be compensated more than one dollar worth of the aggregate consumption commodity in order to give up one dollar worth of health status. The monetary value of the median level of the latent health shocks (θ) is found to be \$3,994 in 1996 dollars.

VI. Examination of Moral Hazard and Adverse Selection

VI.A.1. Moral Hazard

We adopt a definition of “moral hazard” similar to one conventionally used in the health economics literature (e.g., Pauly 1968). In contract theory, instead, this term is reserved for situations in which agent’s behavior cannot be directly observed by the agent, i.e., “hidden actions.” Our definition of moral hazard is closer to the concept of sequential contracting in contract theory, i.e., the agent receives private information in multiple stages during the contractual relationship. Hence the principal first offers a menu

of contracts and the agent chooses one contract. After receiving additional private information the agent selects his best action within the previously chosen contract (see e.g., Courty and Li 2000, Dai, Lewis and Lopomo forthcoming).

We develop and compute a measure of moral hazard that is more general than the one traditionally found in the empirical literature, e.g., Manning et al (1987).¹⁴ Using the structural estimates of the utility parameters and the distribution of latent health status, we compute the elasticity of total medical expenditure with respect to a local change in the rate of co-payment of medical expenditures,

$$\frac{\delta m}{\delta z'(m)} \times \frac{z'(m)}{m}.$$

This represents a counterfactual policy experiment where for each individual there is a marginal change in the co-payment rate with a corresponding change in the demand for medical care. If the co-payment rate was a constant r , then the co-payment schedule would be a linear function of the total expenditure, i.e., $z(m) = r.m$, and our measure of moral hazard would be identical to that used traditionally in the literature (e.g., Pauly 1968).

One advantage of this new measure of moral hazard is that it allows for considerable nonlinearities in the change in behaviors with respect to changes in the reimbursement policy. This allows us to compute a distribution of elasticities for our sample, rather than computing a single statistic of this measure. We compute elasticities for every individual $i = 1, \dots, n$ in our data set using their observed level of total medical expenditures and their associated out of pocket costs, given the estimates of the model parameters. This elasticity is calculated by applying the implicit function theorem to the

¹⁴ Vera-Hernandez (2003) also computes a variant of the traditional measure of moral hazard in the medical care context.

first order condition (2) (or equation 3) and assuming $z(m_i) = m_i \cdot z'(m_i) + a$, where a is a constant which yields,

$$\frac{\delta m_i}{\delta z'(m_i)} \times \frac{z'(m_i)}{m_i} = \frac{-\gamma_1 c_i^{-\gamma_1-1} m_i z'(m_i) - c_i^{-\gamma_1}}{\gamma_1 c_i^{-\gamma_1-1} z'(m_i)^2 + \gamma_2 \gamma_3 (m_i - \theta_i)^{-\gamma_3-1}} \times \frac{z'(m_i)}{m_i}.$$

Table 4 provides the summary statistics of this elasticity conditional on the insurance category. We find that overall at the median, a one percent increase in the rate of co-payment leads to a 0.21 percent drop in total medical expenditures. This number is consistent with that values found in the RAND study (Manning et al. 1987, Newhouse 1993). In particular, Manning et al (1987, pp. 267-268) find that the mean coinsurance elasticity calculated using episodes of treatment for all types of care lies in the range of -0.17 and -0.22. They report that the average coinsurance elasticity calculated using an indirect utility function over medical expenditures (which is conceptually the closest to our calculation) yields an estimate of -0.18. They also report mean values in the range of -0.1 and -0.14 using a calculation of coinsurance elasticity based on average coinsurance rates. That the median of the estimated distribution of elasticities does not differ substantially from that in the existing literature gives us confidence that our model does not suffer from severe misspecification bias.

Our estimation strategy allows us to examine the entire distribution of elasticities which may exhibit a wide range due to nonlinearities in the contracts, i.e., the co-payment schedule. Economic theory predicts (e.g., Pauly 1968) individuals with smaller (and presumably more discretionary) medical expenditures would tend to be more elastic as opposed to those with larger (and presumably more non-discretionary) medical expenditures. We find evidence to substantiate this, i.e., the overall elasticity at the 25th percentile of medical expenditures is -0.47, while at the 75th percentile it is -0.010. This

pattern of variation in elasticities is similar across the different insurance categories. Table 4 also shows that at the median, individuals in the privately purchased insurance category are the least elastic. One potential explanation for this is that there may be self selection of relatively unhealthy individuals in to this insurance category. On the other hand, at the median, individuals in the self-employed insurance category are the most elastic, which may be due to these individuals being relatively healthy. We explore the issue of selection further in section VI.C.

VI.A.2. A Distribution Free Test For Moral Hazard

We propose a distribution free test for moral hazard using our estimates. A prediction of many models of insurance markets is that different contracts exist in the market in order to give consumers choice in economizing on health expenditures. An implication of such theories is that we should expect to see the co-payment elasticity vary across health plans.

We test formally for the differences in the distribution of elasticities across the insurance categories using Kolmogorov-Smirnov test statistics. These are reported in Table 5 and show that in general there is not much statistical difference in the distribution of elasticities across plans. The exceptions are that the distribution of elasticities in the no insurance category is significantly different from those in the employer provided and Medicare categories, and that in employer provided category is significantly different from that in the privately purchased insurance category. In general, the implication is that *all* the insurance categories offer contracts that are designed to induce large elasticities at low expenditures and small elasticities at high expenditures. This may be the consequence of a *common* approach to dealing with moral hazard in these insurance plans. An important implication for empirical research is that it is as important to look at

within plan variation in the elasticity of medical utilization with respect to co-payment as it is to look at variation across plans.

Figure 3 plots the densities and Figure 4 the CDF of the elasticities by insurance category. The density plots look very similar across insurance categories, much like the elasticities reported in Table 4. In each of these the left tail is quite thick as expected with a mode close to zero. In particular Figure 4 shows that the entire distribution of elasticities in the self-employed insurance category is stochastically dominated by the CDFs of the other insurance categories. This suggests that individuals in the self-employed insurance category are the most elastic irrespective of where they lie in the distribution in that category.

VI.A.3. The Relationship Between Observables and Moral Hazard

In order to better understand the factors affecting moral hazard, we examine the correlation between the estimated elasticities and various individual characteristics. The results are reported in Table 6. Our preferred specification is in column IV. We find that individuals who are older and white are more inelastic, where as individuals with higher income are more elastic. Individuals who are in the employer provided, self employed and privately purchased insurance categories are more inelastic relative to those on Medicare, which is the omitted insurance category. Elasticity is monotonically increasing in self reported health status, i.e., individuals who have a higher self reported health status are more elastic (the omitted category is “poor” health). We also find that education is not correlated with elasticity. These results suggest that individuals who are younger and healthier would be more elastic and more responsive to incentives of contracts, and hence a more appealing pool of clients for insurers.

VI.B.1. Adverse Selection.

The results of our estimation procedure allow us to recover measures of the latent health shock, θ , for each individual (using equation 5). We obtain a value for θ in 1996 dollars and a higher value signifies poorer health. To our knowledge we are the first to quantify the latent type of an individual in the context of a model of (health insurance) contracts. Table 7 reports the distribution of the overall latent health status and conditional on insurance category. The median level of overall health shock is about \$4,063 (over two years). Individuals also face substantial risk, as measured by the inter-quartile range of \$10,945, which is more than twice the median value of the health shock.

In comparing across insurance categories, we find that individuals who are not insured have the best latent health status at the median. Similarly, the risk faced by these individuals as measured by the inter-quartile range is the lowest, i.e., \$8,857. This suggests that individuals in good health tend not to purchase insurance due to the lower risk that they face. On the other hand individuals on the VA/Champus plans have the worst latent health status at the median and also face the highest risk with an inter-quartile range of \$18,014. The 75th percentile of the distribution of latent health shocks in this category is much larger than for any of the other categories suggesting that there are some very sick individuals on the VA/Champus plans. This is not surprising since these are veterans who are older than 50 years of age. In general we find that there is great variation in latent health status within each insurance category.

VI.B.2. A Distribution Free Test For Adverse Selection

We propose a distribution free test for adverse selection similar to that for moral hazard. An implication of adverse selection would be that the distribution of the latent health shock varies across health plans. We test for this formally using Kolmogorov-

Smirnov test statistics, which are reported in Table 8. Although, Table 7 suggests that individuals who are not insured are healthier and face a lower risk than others, while those in the VA/Champus plan are unhealthier and face greater risks than others, we cannot reject the equality of the distributions of latent health status across insurance categories. Figure 5 shows the densities of the latent health shock conditional on insurance category. It is seen that all the distributions are skewed to the right and have a probability mass close to zero. Thus a large number of individuals in each insurance category do not suffer large health shocks, but there are some individuals with very large health shocks in each category. The density for VA/Champus has a particularly thick right tail indicating that a sizeable portion of individuals in this category suffer from poor health. Figure 6 shows the CDFs associated with the latent health distribution in each insurance category. Though the CDF of the health shocks in the VA/Champus category is not stochastically dominated by the other CDFs; for a non trivial portion of the support it lies below the other CDFs. This is again evidence that individuals in this category are in poorer health compared to those in other insurance categories.

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, e.g., across different kinds of employer provided plans. Yet, Cardon and Hendel (2001) performed such an analysis for various categories of employer provided plans using the National Medical Expenditure data and found no evidence for

adverse selection in such plans.¹⁵ Thus we view our findings as surprising but not implausible.

VI.B.3. The Relationship Between Observables and Latent Health Status

In order to better understand the nature of asymmetric information we examine the relationship between the latent health shock and observable individual characteristics. The results are reported in Table 9 and our preferred specification is column III. We find that individuals who are older and white experience larger health shocks, and that the self reported health status (“poor” is the omitted category) is positively correlated with latent health status recovered from the estimated model. This is particularly important because our model places no restrictions on the relationship between self reported health status and the latent health shock. Hence this is further evidence that our model captures the trends in the data well and that the estimates are credible. The relationship between education and latent health appears to be non linear but is not significant.¹⁶ We also find that the correlation between observable individual characteristics and the latent health shocks, explains only 5 percent of the variation in the latent health status suggesting that it may not be easy for firms to assess the latent health status of individuals using the typically observed characteristics of individuals. This could give rise to substantial asymmetric information in health insurance markets.

We also we find that there is a positive correlation between the latent health shocks and elasticities (Table 10). This is consistent with a scenario where individuals who have a larger latent health shock (i.e., are more sick) are less elastic. Moreover, in

¹⁵ The lack of detail in our nationally representative HRS data prevent us from doing such an analysis but an advantage of our data are that our results are more generalizable. A disadvantage of using more detailed data that came say from a single firm would be that the results of such a study may not be generalizable.

¹⁶ There is substantial evidence that education and health status are correlated (see e.g., Lleras-Muney 2005). Our finding that education is not correlated with latent health *conditional* on self reported health does not contradict this evidence.

Figure 7 it is seen that the relationship between latent health shocks and elasticities is also highly non linear.

VI.C. Discussion

Our results suggest that the demand for health care is not easily captured by standard models of the insurance market with a separating equilibrium. In these models, a menu of contracts is offered and consumers sort between alternative contracts based on their private information about their health status. The contracts generate an inefficient allocation of resources because insurers inefficiently ration health care by capping benefits for the healthiest consumers.

Our results suggest that sorting between contracts is not particularly well explained by an agent's latent health status. Interestingly, the summary statistics in Table 1 are also fairly consistent with this finding in that the differences in self reported health status do not vary much between different type of health plans. It is difficult to reconcile these estimates with the conventional wisdom about distortions that occur from "cherry picking" in health insurance markets. This is surprising given the importance that adverse selection has received in theoretical models of insurance markets. A partial explanation for this finding is suggested by Figures 1 and 2. Health insurance is most important to consumers when they are very sick and total health expenditures are likely to be high. However, when this is the case, insured consumers on average should expect to be reimbursed for a similar (and large) proportion of their medical expenses irrespective of plan choice. Therefore, the incentive to sort between alternative insurance plans seems limited.

Furthermore, we find that much of the variation in the elasticity of demand for health care occurs within a plan, rather than between plans. Consumers who consume the

least in health service, within a plan, are the most elastic with respect to co-payment rates. Relatively little of the variation in this elasticity is explained by sorting across plans, with the exception of those with no insurance. This seems contrary to models in which a main purpose of multiple plans is to encourage consumers to economize on health expenditures. Our results suggest that an important question for future theoretical research is to explain the nonlinearity in co-payment schedules.

VII. Conclusion

We specify a model of demand for health insurance and medical utilization in the presence of unobserved heterogeneity in the latent health status of individuals. Using a semiparametric procedure we estimate the structural parameters of this model accounting adverse selection and moral hazard due to asymmetric information. We use the estimates of the model to examine the nature of adverse selection and moral hazard in health insurance contracts. We find evidence of moral hazard in health insurance plans but the evidence for adverse selection seems to be lacking. We note that our findings do not support a model of separating equilibrium in the insurance market and our estimates provide a partial explanation for these findings.

Although our proposed semi-parametric method provides a more flexible and robust alternative for analyzing the empirical issues of adverse selection and moral hazard in health insurance, several limitations are acknowledged. The utility function specification we use is assumed to be separable in the consumption of the composite good and health status. While this specification captures the risk aversion features of consumer utilities in health status, it rules out more flexible interactions between the utility derived from composite good consumption and health status. It is sometimes argued that the marginal utility for composite good consumption might decrease in the

case of severe illness (e.g. Viscusi and Evans 1990). We note, however, that Spence and Zeckhauser (1971) and Blomqvist (1997) use a similar specification, and Campo, Guerre, Perrigne and Vuong (2003) also require similar restrictions on utility in an auctions context. The utility specification also assumes that medical utilization is a perfect substitute for an individual's health status. This is similar to the assumption made by Cardon and Hendel (2001). However, we do allow for a flexible nonlinear relationship between medical utilization and latent health status, which we expect does incorporate some of the stochasticity in the health production relationship. Due to data limitations we do not allow for heterogeneity in risk preferences, which can be a potentially important determinant of moral hazard and selection in to contracts. However in principle our methodology can be extended to incorporate heterogeneity in risk preferences with access to better data.

In conclusion, in spite of these limitations 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. Our research is also important as it provides a framework for similar analysis in other contexts, especially with cross section data, where distortions exist due to asymmetric information.

References:

- Abbring, J. H., P. A. Chiappori, J. H. Heckman and J. Pinquet. Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish? *Journal of the European Economic Association*. 2003; 1(Papers and Proceedings): 512-521.
- Akerlof, George. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*. 1970; 84(3): 488-500.
- Arrow, K.J. Uncertainty and the Welfare Economics of Medical Care. *American Economic Review*. 1963; 53:941-973.
- Athey, S. and P. Haile. Identification of Standard Auction Models. *Econometrica*. 2002; 70:2107-2140.
- Bajari, Patrick and Matt Kahn. Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities. *Journal of Business and Economic Statistics*. 2005; 23(1): 20-33.
- Blau, D. and Donna B. Gilleskie. The Role of Retiree Health Insurance in the Employment Behavior of Older Males. Working Paper. UNC. 2003.
- Blomqvist, A.G. Optimal Non-linear Health Insurance. *Journal of Health Economics*. 1997; 16(3):303-321.
- Cameron, A. C.; P. K. Trivedi, Frank Milne, and J. Piggott. A Microeconometric Model of the Demand for Health Care and Health Insurance in Australia. *Review of Economic Studies*. 1988; 55(1):85-106.
- Campo, Sandra M.; I. Perrigne, and Q. Vuong. Asymmetry and Joint Bidding in OCS Wildcat Auctions. *Journal of Applied Econometrics*. 2003; 18:179-207.
- Campo, Sandra M.; E. Guerre, I. Perrigne, and Q. Vuong. Semiparametric Estimation of First-Price Auctions with Risk Averse Bidders. Working Paper. UNC. 2003.
- Cardon, James H. and Igal Hendel. Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey. *RAND Journal of Economics*. 2001; 32(3):408-27.
- Cawley, J., and T. Philipson. An Empirical Examination of Information Barriers to Trade in Insurance. *American Economic Review*, 1999; 89: 827-46.
- Chiappori, Pierre-Andre and B. Salanié. Testing for Asymmetric Information on Insurance Markets. *Journal of Political Economy*. 2000; 108(1):56-78.

- Chiappori, Pierre-Andre and Bernard Salanié. Testing Contract Theory: a Survey of Some Recent Work, in *Advances in Economics and Econometrics – Theory and Applications*, Eighth World Congress, M. Dewatripont, L. Hansen and P. Turnovsky, ed., *Econometric Society Monographs*, Cambridge University Press, Cambridge, 2003: 115-149.
- Chiappori, Pierre-Andre, Bruno Jullien, and François Salanié. *Asymmetric Information in Insurance: General Testable Implications*. *RAND Journal of Economics*. Forthcoming.
- Cohen, Alma, and Liran Einav. *Estimating Risk Preferences from Deductible Choice*. NBER Working Paper No. 11461. 2005.
- Courty, P., and Hao Li. *Sequential Screening*. *Review of Economic Studies*. 2000; 67:697-717.
- Currie, Janet and Brigitte C. Madrian. *Health, Health Insurance and the Labor Market*. Ashenfelter, Orley and David Card (eds) *Handbook of Labor Economics* Vol. 3C. Amsterdam ; New York : Elsevier; 1999: 3309-3416.
- Cutler, David M. and Richard J. Zeckhauser. *The Anatomy of Health Insurance*. Culyer, Anthony J. and Joseph P. Newhouse (eds) *Handbook of Health Economics* Vol. 1A. Amsterdam ; New York : Elsevier; 2000: 563-643.
- Dai, Chifeng, Tracy Lewis, and Giuseppe Lopomo. *Delegating Management to Experts*. *RAND Journal of Economics*. Forthcoming.
- De Meza, David, and David Webb. *Advantageous Selection in Insurance Markets*. *RAND Journal of Economics*. 2001; 32(2):249-262.
- Fan, J. and Gijbels, I. *Local Polynomial Modelling and Its Applications: Monographs on Statistics and Applied Probability*. London: Chapman & Hall. 1996.
- Ferrall, Christopher and Bruce Shearer. *Incentives and Transactions Costs Within the Firm: Estimating an Agency Model Using Payroll Records*. *Review of Economic Studies*. 1999; 66:309-338.
- Finkelstein, Amy and James Poterba. *Adverse Selection in Insurance Market: Policyholder Evidence from the U.K. Annuity Market*. *Journal of Political Economy*. 2004; 112(1):183-208.
- Finkelstein, Amy and James Poterba. *Testing for Adverse Selection with “Unused Observables.”* NBER Working Paper No. 12112. 2006.

- Finkelstein, Amy and Kathleen McGarry. Multiple Dimensions for Private Information: Evidence from the Long-Term Care Insurance Market. *American Economic Review*. Forthcoming.
- Gertler, Paul and Jonathan Gruber. Insuring Consumption Against Illness. *American Economic Review*. 2002; 92(1): 51-70.
- Gilleskie, Donna B. A Dynamic Stochastic Model of Medical Care Use and Work Absence. *Econometrica*. 1998; 66(1):1-45.
- Gourinchas, P. O. and J. Parker. Consumption Over the Life Cycle. *Econometrica*. 2002. 70(1): 47-89.
- Gruber, Jonathan. Health Insurance and the Labor Market. Culyer, Anthony J. and Joseph P. Newhouse (eds) *Handbook of Health Economics Vol. 1A*. Amsterdam ; New York : Elsevier; 2000; pp. 645-706.
- Guerre, E.; I. Perrigne, and Q. Vuong. Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica*. 2000; 68(3):525-574.
- Hansen, L. P. and K. J. Singleton. Generalized Instrumental Variables Estimation of Non Linear Rational Expectations Models. *Econometrica*. 1982; 50(5):1269-1285
- Harris, K. and M. Keane. A Model of Health Plan Choice: Inferring Preferences and Perceptions From a Combination of Revealed Preferences and Attitudinal Data. *Journal of Econometrics*. 1999; 89: 131-157.
- Heckman, James. J. and Thomas E. MaCurdy. A Life Cycle Model of Female Labor Supply. *Review of Economic Studies*. 1980; 47(2): 47-74.
- Hong, Han and Elie Tamer. Inference in Censored Models with Endogenous Regressors. *Econometrica*. 2003; 71(3):905-932.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. Precautionary Savings and Social Insurance. *Journal of Political Economy*. 1995; 103(2): 360–99.
- Juster, F.T. and R. Suzman. The Health and Retirement Study: An Overview. *Journal of Human Resources*. 1995; Supplement (JHR 30-S): S7-S56.
- Keeler, E., J. Newhouse and C. Phelps. Deductibles and the Demand for Medical Care Services: The Theory of a Consumer Facing a Variable Price Schedule under Uncertainty. *Econometrica* 1977; 45: 641-655.
- Khwaja, Ahmed. Health Insurance Habits and Health Outcomes: A Dynamic Stochastic Model of Investment in Health. Ph.D Dissertation, University of Minnesota. 2001.

- . The Dynamic Trade-off Between Medical Expenditure Insurance, Moral Hazard and Mortality Risk: A Life Cycle Analysis of the Effects of Medicare on Individual Health Incentives and Health Outcomes. Working Paper. Duke University, Fuqua School of Business; 2005.
- Lleras-Muney, A. The Relationship Between Education and Adult Mortality in the U.S. *Review of Economic Studies*. 2005; 72(1): 189-221.
- Manning, W.G.; J.P. Newhouse, N.Duan et al. Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment. *American Economic Review*. 1987; 77(3):251-277.
- Newey, W. K. The Asymptotic Variance of Semiparametric Estimators. *Econometrica*. 1994; 62(6):1349-1382.
- Newhouse, J.P. *Free For All? Lessons from the Health Insurance Experiment*. Cambridge: Harvard University Press. 1993.
- Paarsch, Harry J. and Bruce Shearer. Piece Rates, Fixed Wages, and Incentive Effects: Statistical Evidence from Payroll Records. *International Economic Review*. 2000; 41(1):59-92.
- Pauly, M. The Economics of Moral Hazard: Comment. *American Economic Review*. 1968; 58:531-536.
- . Overinsurance and Public Provision of Insurance: the Roles of Moral Hazard and Adverse Selection. *Quarterly Journal of Economics*. 1974; 88(1):44-54.
- . Insurance Reimbursement. Culyer, Anthony J. and Joseph P. Newhouse (eds) *Handbook of Health Economics Vol. 1A*. Amsterdam ; New York : Elsevier; 2000; pp. 537-560.
- Powell, James. Least Absolute Deviations Estimation for the Censored Regression Model. *Journal of Econometrics*. 1984; 25:303-325.
- Rothschild, M. and J.E. Stiglitz. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics*. 1976; 90(4):630-649.
- Sieg, Holger. Estimating a Dynamic Model of Household Choices in the Presence of Income Taxation. *International Economic Review*. 2000; 41 (3), 637-668.
- Shea, J. Union Contracts and the Life-Cycle/Permanent-Income Hypothesis. *American Economic Review*. 1995; 85(1): 186-200.

- Spence, M. Job Market Signaling. *The Quarterly Journal of Economics*. 1973; 87(3): 355-374.
- Spence, M. and R. Zeckhauser. Insurance, Information, and Individual Action. *American Economic Review*. 1971; 61(2):380-387.
- Viscusi, K. and W. Evans. Utility Functions that Depend on Health Status: Estimates and Economic Implications. *American Economic Review*. 1990; 80: 353-374.
- Wilson, C. A Model of Insurance Markets with Incomplete Information. *Journal of Economic Theory*. 1977; 16:167-207.
- Wilson, C. The Nature of Equilibrium in Markets With Adverse Selection. *Bell Journal of Economics*. 1980; 11(1):108-130.
- Vera-Hernandez, Marcos. Estimation of a Principal-Agent Model: Moral Hazard in Medical Insurance. *The Rand Journal of Economics*. 2003. 34(4):670-693.
- Zeldes, S. P. Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*. 1989; 97(2):305-346.
- Zeckhauser, R. Medical Insurance: a Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives. *Journal of Economic Theory*. 1970; 2(1):10-26.
- Zweifel, Peter and Willard G. Manning. Moral Hazard and Consumer Incentives in Health Care. Culyer, Anthony J. and Joseph P. Newhouse (eds) *Handbook of Health Economics* Vol. 1A. Amsterdam ; New York : Elsevier; 2000; pp. 409-459.

Appendix I: Variable Descriptions

Variable	Variable Description
<i>HRS Data</i>	
Health status	Self reported health status: 6-excellent, 5-very good, 4-good, 3-fair, 2 – poor
HH size	Number of members in the household
Insurance categories	No insurance, Employer provided insurance, Self employed with insurance through own business, VA/Champus, Privately purchased insurance, Medicare
HH OOP med. exp.	total household out of pocket (OOP) cost for med. Care
HH tot. med. exp.	total imputed household medical exp. (sum of OOP + insurer payment)
Rural	<p>A measure (0-9) of how rural an area is.</p> <p>0. Central counties of metropolitan areas of 1 million population or more</p> <p>1. Fringe counties of metropolitan areas of 1 million population or more</p> <p>2. Counties in metropolitan areas of 250 thousand to 1 million population</p> <p>3. Counties in metropolitan areas of than 250 thousand population</p> <p>4. Urban population of 20,000 or more, adjacent to a metropolitan area</p> <p>5. Urban population of 20,000 or more, not adjacent to a metropolitan area</p> <p>6. Urban population of 2,500 to 19,999, adjacent to a metropolitan area</p> <p>7. Urban population of 2,500 to 19,999, not adjacent to a metropolitan area</p> <p>8. Completely rural or less than 2,500 urban population, adjacent to a metropolitan area</p> <p>9. Completely rural or less than 2, urban population, not adjacent to a metropolitan area</p>
Urban	<p>1 if Rural = 0</p> <p>0 otherwise</p>
<i>Non HRS Data</i>	
Annual payroll	<p>Annual Payroll-available at County level</p> <p>Source: US Census Bureau (County Business Patterns)</p> <p>http://www.census.gov/epcd/cbp/view/cbpview.html</p> <p>Total payroll includes all forms of compensation, such as salaries, wages, reported tips, commissions, bonuses, vacation allowances, sick-leave pay, employee contributions to qualified pension plans, and the value of taxable fringe benefits. For corporations, it includes amounts paid to officers and executives; for unincorporated businesses, it does not include profit or other compensation of proprietors or partners. Payroll is reported before deductions for Social Security, income tax, insurance, union dues, etc. First-quarter payroll consists of payroll during the January-to-March quarter.</p>
Employees	<p>Employees per week-available at County level</p> <p>Source : US Census Bureau (County Business Patterns)</p> <p>http://www.census.gov/epcd/cbp/view/cbpview.html</p> <p>Paid employment consists of full- and part-time employees, including salaried officers and executives of corporations, who are on the payroll in the pay period including March 12. Included are employees on paid sick leave, holidays, and vacations; not included are proprietors and partners of unincorporated businesses.</p>
Establishments	<p>Total Establishments –available at County level</p> <p>Source : US Census Bureau (County Business Patterns)</p> <p>http://www.census.gov/epcd/cbp/view/cbpview.html</p> <p>An establishment is a single physical location at which business is conducted or services or industrial operations are performed. It is not</p>

	<p>necessarily identical with a company or enterprise, which may consist of one or more establishments. When two or more activities are carried on at a single location under a single ownership, all activities generally are grouped together as a single establishment. The entire establishment is classified on the basis of its major activity and all data are included in that classification. Establishment-size designations are determined by paid employment in the mid-March pay period. The size group "1 to 4" includes establishments that did not report any paid employees in the mid-March pay period but paid wages to at least one employee at some time during the year.</p> <p>Establishment counts represent the number of locations with paid employees any time during the year. This series excludes governmental establishments except for wholesale liquor establishments (NAICS 4228), retail liquor stores (NAICS 44531), Federally-chartered savings institutions (NAICS 522120), Federally-chartered credit unions (NAICS 522130), and hospitals (NAICS 622).</p>
HPI	<p>Housing Price Index—Available at State level Source : Office of Federal Housing Enterprise Oversight http://www.ofheo.gov/HPI.asp HPI is a measure designed to capture changes in the value of single family homes. HPI is a weighted repeat sales index that measures average price changes in repeat sales or refinancings on the same properties. Data provided by Fannie Mae and Freddie Mac.</p>
Insurance coverage	<p>People covered by insurance- Available at State level. Source : US Census Bureau CPS 1997 http://www.census.gov/hhes/www/hlthins/cover96/c96tabf.html Includes private and government sponsored plans</p>
GPCI (Work, Malpractice, Practice components)	<p>Available at county level. Source : Medicare Payment Advisory Commission http://www.mgma.com/research/gpci.cfm A Geographic Practice Cost Index (GPCI) is used by Medicare to adjust for variance in operating costs of medical practices located in different parts of the country. Reimbursement of Physicians for services performed under Medicare is governed by a formula that considers the product of three factors:</p> <ol style="list-style-type: none"> 1. A nationally uniform relative value unit (RVU) for the service; 2. A GPCI value which adjusts each RVU component (Work, Practice Expense, malpractice); 3. A nationally uniform conversion factor for the service. <p>The Conversion Factor converts the relative values into payment amounts. For each physician fee schedule service, which is represented by an associated Health Care Common Procedure Coding System (HCPCS) code, there are three relative values:</p> <ol style="list-style-type: none"> 1. An RVU for physician work; 2. An RVU for practice expense; 3. An RVU for malpractice expense. <p>For each of these components, there is a GPCI which adjusts the RVU value based on a practices geographic location. The GPCIs reflect the relative costs of practice expenses, malpractice insurance, and physician work in an area compared to the national average for each component.</p>
Unemployment Rate	<p>Available at county level. Source : U.S. Census Bureau http://censtats.census.gov/usa/usa.shtml</p>

Table 1: Summary Statistics

Variable Description	Obs	Mean	Std. Dev.
<i>HRS Data</i>			
Health status			
All insurance categories	3648	4.49	1.07
No Insurance	316	3.93	1.15
Employer Provided	2644	4.56	1.03
Self Employed	81	4.93	.891
VA/Champus	123	4.33	1.18
Privately Purchased	255	4.67	1.00
Medicare	229	4.14	1.16
Household size	3648	2.40	1.14
Household income	3648	53917.66	47892.64
Insurance category	3648	2.46	1.29
Hh insurance premium	3648	1363.08	2178.08
Total household OOP cost	3648	1758.87	2622.31
Total medical expenses	3648	11122.24	13237.82
Age	3648	58.6187	5.64
Urban	3648	0.41	0.49
<i>Non HRS Data</i>			
Annual payroll county	3648	1.01e+07	1.90e+07
# employees county	3648	328441.3	574392.8
#establishments county	3648	19917.4	33946.61
Housing price index	3648	195.74	36.56
People covered by ins in state	3648	9051.52	6802.55
Geog Practice Cost index (work)	3648	.990	.026
Geog Practice Cost index (malpractice)	3648	.944	.458
Geog Practice Cost index (practice)	3648	.961	.096
Unemployment Rate	3648	5.44	2.61

Table 2: Summary Statistics of Local Linear Regression

Variable	Obs	Mean	Std. Dev.	Min	Max
Conditional Mean (α)	3648	1676.14	1131.57	450.46	15079
Medical expenditure	3648	.109	.033	.0004	.369
No insurance	3648	-342.58	429.13	-881.14	2901
Employer provided	3648	-543.38	285.49	-1343.8	2350.2
Self employed	3648	216.27	1327.95	-2497.6	11107
VA/Champus	3648	-762.09	718.23	-1631.5	7496.2
Medicare	3648	-579.36	477.26	-2287.3	4382.9
Urban	3648	-56.19	333.39	-2771.4	8206.5

The above table gives the results from a local linear regression estimation of the supply side of the model
The dependent variable is *OOP Medical Expenses*. Privately purchased insurance was the omitted insurance category.

Each individual's α can be interpreted as their expected *OOP Medical Expenses* conditional on their observable characteristics.

Each individual's coefficient on the covariates can be interpreted as the derivative of this conditional expectation with respect to the covariate.

Table 3: Estimates of Stuctural Utility Paramters

Coefficient of Relative Risk Aversion on Consumption (γ_1)	Relative Utility Weight for Health (γ_2)	Coefficient of Relative Risk Aversion on Health (γ_3)	Median Health Shock (θ)	Objective Function
0.846 (0.1278)	1.371 (0.592)	1.521 (0.203)	3994.33 (1982.71)	0.827

Standard errors are in parentheses

Table 4: Elasticity of Total Medical Expenditure w.r.t. the Effective Price of Medical Expenditure by Insurance Category

Insurance Category	25 th percentile	Median	75 th Percentile
All Insurance Categories	-0.47	-0.21	-0.01
No Insurance	-0.44	-0.18	-0.06
Employer Provided	-0.47	-0.21	-0.1
Self Employed	-0.64	-0.24	-0.12
VA/Champus	-0.48	-0.19	-0.07
Privately Purchased	-0.42	-0.16	-0.07
Medicare	-0.45	-0.23	-0.11

Elasticity of medical expenditure wrt the effective price of medical expenditure is calculated as:

$$\frac{\delta m}{\delta z'(m)} \times \frac{z'(m)}{m}$$

Applying implicit function rule to First Order Condition and assuming $z(m) = m \cdot z'(m) + cons$ yields:

$$\frac{\delta m}{\delta z'(m)} \times \frac{z'(m)}{m} = \frac{-\gamma_1 c^{-\gamma_1-1} m z'(m) - c^{-\gamma_1}}{\gamma_1 c^{-\gamma_1-1} z'(m)^2 + \gamma_2 \gamma_3 (m - \theta)^{-\gamma_3-1}} \times \frac{z'(m)}{m}$$

Table 5: Kolmogorov-Smirnov Tests for Equality of Distributions of Price Elasticities

Insurance Category	Employer Provided	Self Employed	VA/Champus	Privately Purchased	Medicare
No Insurance	0.175 (0.000)	0.168 (0.325)	0.112 (0.95)	0.106 (0.35)	0.194 (0.000)
Employer Provided		0.112 (0.60)	0.125 (0.675)	0.149 (0.000)	0.039 (1.0)
Self Employed			0.146 (0.725)	0.18 (0.125)	0.132 (0.70)
VA/Champus				0.097 (0.925)	0.116 (0.925)
Privately Purchased					0.1735 (0.125)

The pairwise Kolmogorov-Smirnov Statistic is reported with corrected P-Values in parentheses

Table 6: OLS Results from Elasticity Regressions

	I Elasticity	II Elasticity	III Elasticity	IV Elasticity
Urban	-0.016 [0.020]	-0.019 [0.020]	-0.007 [0.020]	-0.007 [0.020]
HH size	0.02 [0.008]*	0.022 [0.008]**	0.02 [0.008]*	0.02 [0.008]*
Age	0.005 [0.002]*	0.007 [0.002]**	0.006 [0.002]*	0.006 [0.002]*
White	0.061 [0.026]*	0.058 [0.026]*	0.087 [0.026]**	0.089 [0.026]**
Male	0.002 [0.022]	0.006 [0.022]	0.003 [0.021]	0.002 [0.021]
HH income	-2.1E-06 [3.0E-07]**	-2.1E-06 [3.0E-07]**	-1.8E-06 [3.0E-07]**	-1.8E-06 [4.0E-07]**
Insurance category				
No Insurance		0.134 [0.057]*	0.107 [0.056]	0.106 [0.056]
Employer Provided		0.139 [0.048]**	0.15 [0.047]**	0.15 [0.047]**
Self Employed		0.199 [0.089]*	0.216 [0.087]*	0.218 [0.088]*
VA/Champus		0.098 [0.071]	0.094 [0.070]	0.094 [0.070]
Privately Purchased		0.173 [0.060]**	0.196 [0.060]**	0.196 [0.060]**
Health status				
Fair			-0.072 [0.028]*	-0.071 [0.028]*
Good			-0.178 [0.028]**	-0.176 [0.029]**
Very good			-0.256 [0.029]**	-0.255 [0.030]**
Excellent			-0.34 [0.037]**	-0.339 [0.037]**
Education				
12 years				-0.005 [0.021]
13-15 years				-0.014 [0.029]
16 years				-0.021 [0.042]
More than 16 years				0.018 [0.038]
Constant	-0.659 [0.130]**	-0.941 [0.161]**	-0.705 [0.160]**	-0.696 [0.161]**
Observations	3648	3648	3648	3648
R-squared	0.03	0.04	0.06	0.06

Robust standard Errors are reported in brackets. * significant at 5%; ** significant at 1%

Table 7: Health Shocks by Insurance Category

Insurance Category	25 th percentile	Median	75 th Percentile
All Insurance Categories	708.08	4063.45	11653
No Insurance	445.73	2648.85	9302.75
Employer Provided	726	4170	11696
Self Employed	88.43	3551.7	10862
VA/Champus	962	5896.6	18976
Privately Purchased	770	4113	12097
Medicare	782	4408.9	10740

Table 8: Kolmogorov-Smirnov Tests for Equality of Distributions of Health Shocks

Insurance Category	Employer Provided	Self Employed	VA/Champus	Privately Purchased	Medicare
No Insurance	0.103 (0.125)	0.093 (0.80)	0.165 (0.375)	0.103 (0.150)	0.112 (0.35)
Employer Provided		0.113 (0.50)	0.116 (0.65)	0.058 (0.65)	0.045 (0.975)
Self Employed			0.15 (0.65)	0.106 (0.755)	0.101 (0.825)
VA/Champus				0.135 (0.55)	0.141 (0.625)
Privately Purchased					0.054 (0.975)

The pairwise Kolmogorov-Smirnov Statistic is reported with corrected P-Values in parentheses

Table 9: OLS Results from Recovered Health Shock Regressions

	I Health Shock	II Health Shock	III Health Shock
Age	130.12 [40.04]**	75.249 [39.72]	80.609 [40.35]*
White	1,060.80 [493.25]*	2,288.25 [489.32]**	2,184.14 [502.63]**
Male	270.034 [432.53]	285.853 [420.92]	305.074 [424.45]
Health status			
Fair		-4,099.55 [1,459.13]**	-4,146.58 [1,457.8]**
Good		-7,808.83 [1,362.70]**	-7,952.54 [1,365.04]**
Very good		-10,062.47 [1,342.37]**	-10,241.81 [1,349.53]**
Excellent		-10,914.32 [1,374.61]**	-11,103.97 [1,383.97]**
Education			
12 years			498.604 [564.40]
13-15 years			1,124.19 [687.91]
16 years			591.61 [853.46]
More than 16 years			143.919 [725.47]
Constant	150.385 [2,322.67]	10,697.90 [2,734.11]**	10,126.68 [2,786.15]**
Observations	3648	3648	3648
R-squared	0.01	0.05	0.05

The omitted health status category is “poor.” The omitted education level is “less than 12 years”.


Robust standard errors are reported in brackets.

** indicates a variable is statistically significant at the 1% level, * indicates significance at the 5% level.

Table 10: Correlations Between Health Shocks and Elasticities by Insurance Category


Insurance Category	Correlation
All Insurance Categories	0.424
No Insurance	0.40
Employer Provided	0.438
Self Employed	0.514
VA/Champus	0.448
Privately Purchased	0.385
Medicare	0.355

Figure 1: Expected OOP Medical Expenses conditional on Total Medical Expenses



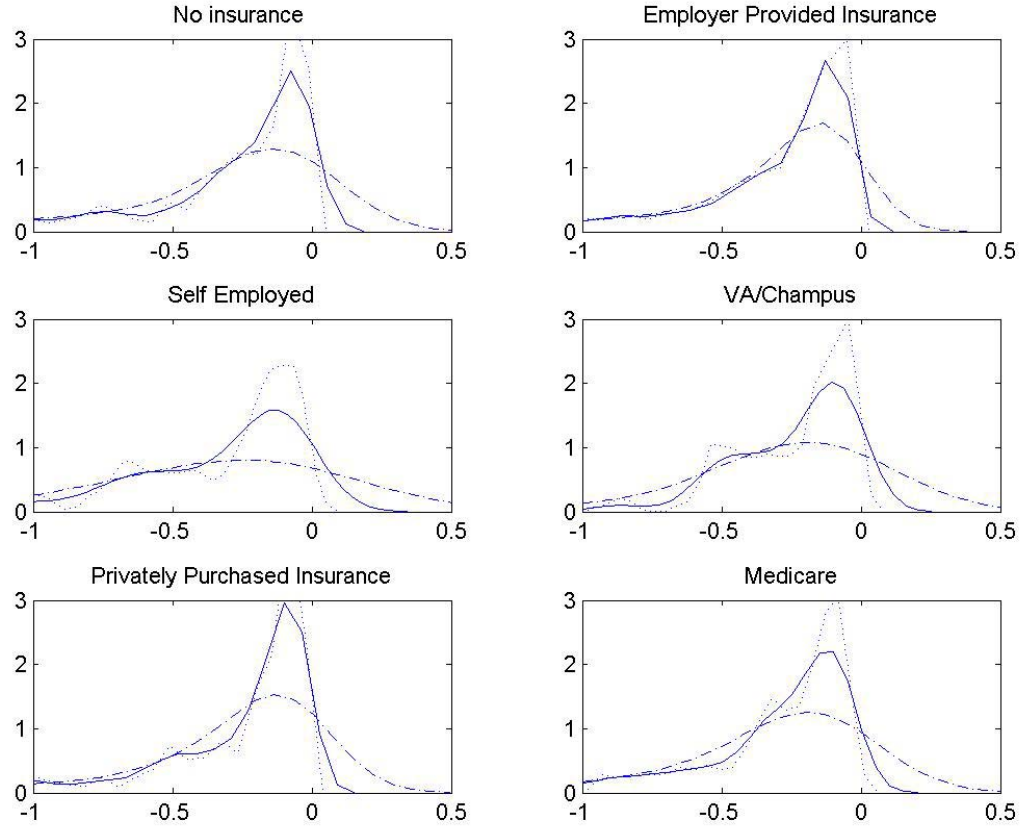
Expected OOP Expenditure is calculated as the fitted values from a local linear regression of *OOP Expenditure* on *Total Medical Expenditure* and a series of dummy variables for insurance categories and rural location. Expected OOP Expenditure is calculated for 50 evenly spaced values of *Total Medical Expenditure* with the other covariates held at their means.

Figure 2: Derivative of OOP Medical Expenses with respect to Total Medical Expenses



The derivative of Expected OOP Expenditure with respect to *Total Medical Expenditure* is calculated from a local linear regression of *OOP Expenditure* on *Total Medical Expenditure* and a series of dummy variables for insurance categories and rural location. The derivative of Expected OOP Expenditure with respect to *Total Medical Expenditure* is calculated for 50 evenly spaced values of *Total Medical Expenditure* with the other covariates held at their means.

Figure 3: Density Estimates for Price Elasticities by Insurance Category



This figure plots estimated densities for the recovered elasticity for each of the six insurance categories. The estimated densities are calculated using a normal kernel. Each subplot estimates the density using three different bandwidths where the default bandwidth is the optimal bandwidth calculated in Matlab. The solid line is the optimal bandwidth, the dashed line is three times the optimal bandwidth, and the dotted line is one third the optimal bandwidth.

Figure 4: CDFs of Elasticity Distributions by Insurance Category


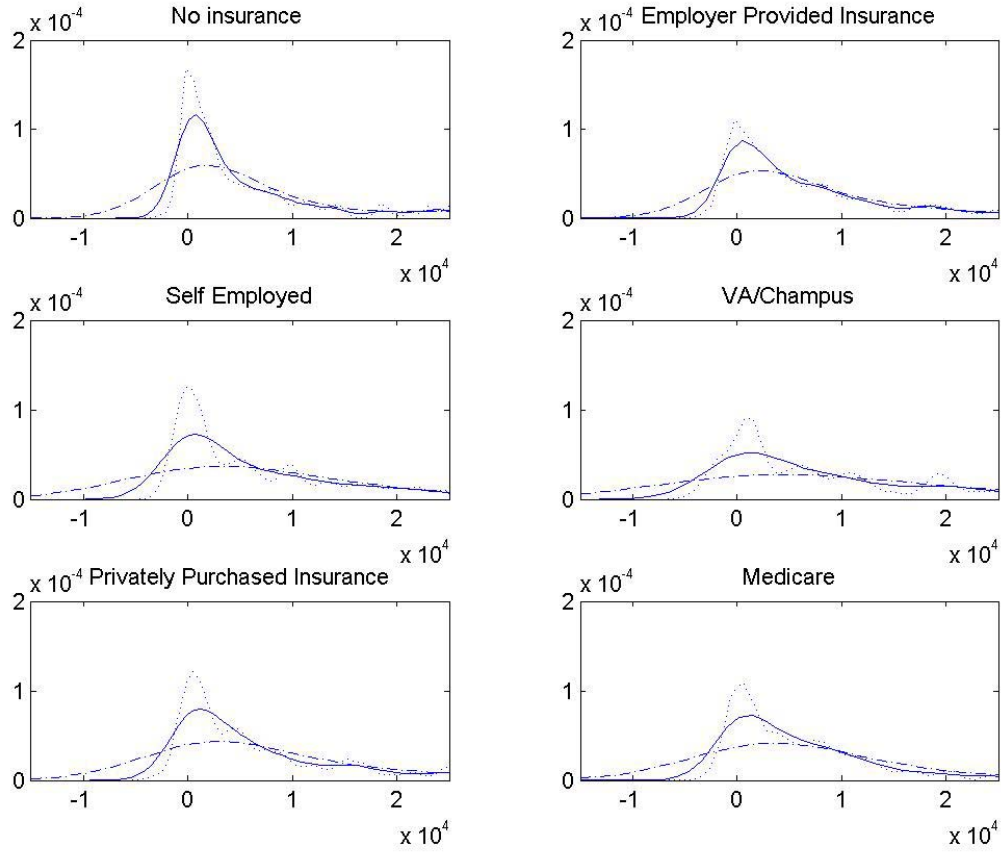


Figure 5: Density Estimates for Recovered Health Shock by Insurance Category



This figure plots estimated densities for the recovered health shock, θ , for each of the six insurance categories.

The estimated densities are calculated using a normal kernel. Each subplot estimates the density using three different bandwidths where the default bandwidth is the optimal bandwidth calculated in Matlab. The solid line is the optimal bandwidth, the dashed line is three times the optimal bandwidth, and the dotted line is one third the optimal bandwidth.

Figure 6: CDFs of Health Shock Distributions by Insurance Category

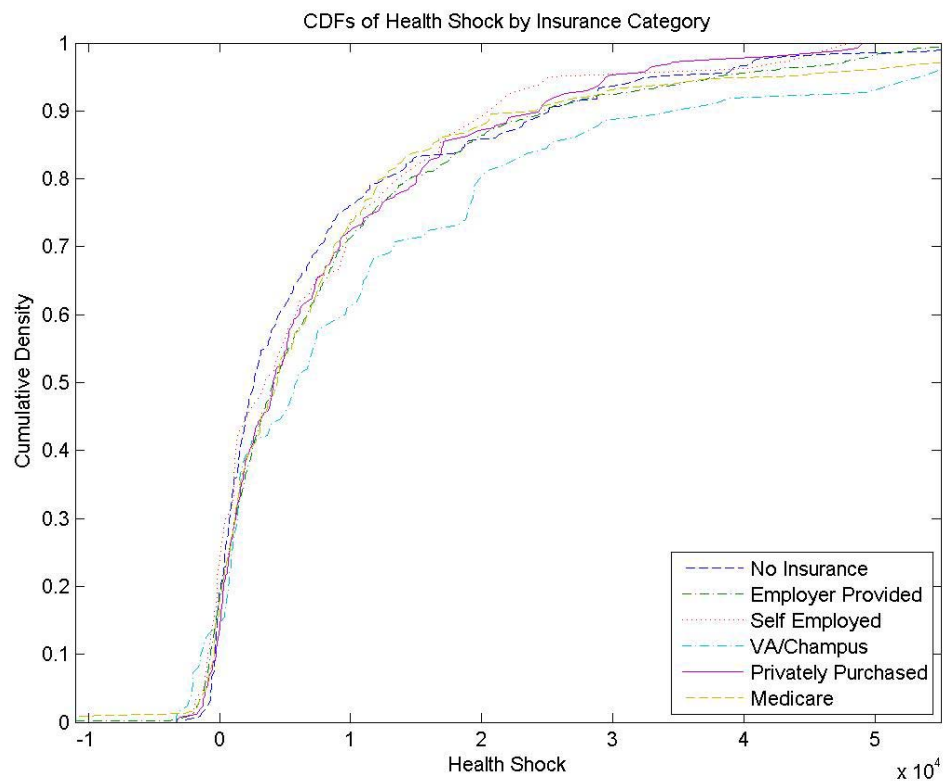


Figure 7: Scatterplots of Elasticities against Health Shocks by Insurance Category

