Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare*

Josh Lustig†

Boston University

Department of Economics

July 2009

Abstract

This paper measures welfare losses caused by adverse selection and imperfect competition in privatized Medicare. I model insurers' premium and coverage choices in an environment where consumers with heterogeneous preferences may impose different costs on their insurers. Using insurers' behavior and variation in market structure, I identify a causal link between consumers' types and insurers' costs and infer whether preferences, which determine insurance demand, contain information about expected health. The results suggest adverse selection is present. Simulated equilibria without adverse selection imply the resulting welfare losses are substantial, 8.1% of total costs, but smaller than those caused by imperfect competition.

I. Introduction

Since Akerlof (1970) and Rothschild and Stiglitz (1976) first formalized the theory of adverse selection, theorists have emphasized the market failures that adverse selection can cause in insurance markets. Subsequent to these theoretical contributions, an empirical literature emerged that tests for adverse selection using consumers' observed insurance choices and their risk outcomes. Recently, methods to detect adverse selection and distinguish it from other information asymmetries have grown increasingly sophisticated. Despite these advances, the current literature is only beginning to attempt to quantify the market failures emphasized in the theory. Market failures from adverse selection arise from distorted consumer and insurer behavior, and thus, measuring them necessitates a model of consumer and insurer behavior. This paper makes two contributions to the adverse selection literature. First, I develop a new strategy to measure adverse selection using the information contained in insurers' coverage and premium choices, and provide evidence of adverse selection in privatized Medicare. Second, and more importantly, I estimate the welfare losses from the resulting market failures.

I estimate a model in which insurers choose how much coverage to offer and what premiums to charge under varying degrees of adverse selection. Using insurers' optimal premium and coverage choices, I calculate insurers' costs of providing insurance coverage. To measure adverse selection, I relate these costs to enrollees' preferences for insurance, which are hidden from insurers but can be inferred using a model of consumer sorting. The exercise enables me to detect the presence of adverse selection. Specifically, do consumers' preferences for insurance, which determine how much health insurance they purchase, contain information about their expected health?

I apply my model to an insurance market where policy makers are increasingly introducing competition: the market for the provision of managed care options to Medicare beneficiaries. Between 2000 and 2003, the Medicare + Choice program (M + C) allowed Medicare beneficiaries to opt out of traditional, fee-for-service Medicare and enroll in plans offered by private insurers.[1] Underlying the M + C program was the belief that competing HMOs could more efficiently provide insurance coverage to the nation's elderly.

The M + C market is one with endogenously differentiated products. HMOs choose how many plans to offer, how much insurance coverage to provide, and what premiums to charge. I assume HMOs' decisions to offer prescription drug insurance and other forms of coverage are made simultaneously in a static setting characterized by a Nash equilibrium. I model consumers' enrollment decisions with a discrete choice framework. Consumer sorting between HMO plans and traditional Medicare depends on a distribution of preferences for insurance as well as consumers' choice sets. The model's structure allows me to infer insurers' costs and to make predictions about consumer sorting across insurance plans. These form the basis for my measure of adverse selection.

To measure adverse selection, I estimate a causal relationship between consumers' preferences (implied by their enrollment decisions) and insurers' costs (implied by their coverage and premium choices). If attracting consumers with strong preferences for insurance coverage *causes* insurers to have high costs, then consumers with strong preferences must tend to have poor expected health. If consumers' preferences depend on characteristics such as risk aversion, in addition to unobserved health, my model could estimate a negative relationship between insurers' costs and consumers' preferences.[2]

My estimation strategy is complicated by unobserved factors that also affect insurers' costs, including moral hazard. HMOs offering generous prescription drug coverage, for example, will attract many consumers with strong preferences for insurance, who may or may not have poor expected health. If moral hazard is present, hidden action will make these consumers appear costly, independent of adverse selection. Similarly, insurers' implied costs become uninformative about adverse selection if insurers' unobserved costs vary and influence coverage choices. Insurers offering generous prescription drug coverage, for example, may have advantageous relationships with drug companies that I am unable to observe.

Thus, my strategy to identify and measure adverse selection exploits a relationship between adverse selection, insurers' costs, and variation in market structure. If and only if adverse selection exists, insurers that enroll many consumers with weak preferences for insurance (and therefore, good expected health) will have lower average costs. Markets with distinct structures provide insurers different opportunities to attract these consumers. In markets with few insurers, for example, HMOs offering generous insurance coverage attract more consumers with weak preferences for generous insurance. In markets with more insurers, however, the product space becomes saturated and these consumers are lost to plans that are less generous and less expensive. I use exogenous variation in market structure and insurers' characteristics to generate observable variation in sorting by consumers with weak and strong preferences for insurance. My measure of adverse selection is identified using only the portions of insurers' costs that are explained by this variation in preference-based sorting.[3]

Estimates of HMOs' cost functions provide evidence of adverse selection. HMOs' total costs of coverage are 7.6% higher for a consumer whose preferences imply a willingness to pay for insurance that is one standard deviation above the median. Market failures arise from this adverse selection because insurers are unable to vary premiums with enrollees' costs. Inefficient pricing leads some consumers to purchase suboptimal levels of insurance to avoid subsidizing the unhealthy.[4] To quantify these distortions, I remove the effects of adverse selection via perfect risk adjustment and type-dependent taxes and subsidies. Generous insurers no longer charge premiums reflecting their unhealthy pools of enrollees and consumers pay premiums that reflect the costs that *they* impose on their insurers but the imperfectly competitive nature of the market is retained. I simulate how insurers' plans and consumers' choices change and find the welfare effects from adverse selection to be substantial; after its removal total surplus increases by up to 8.1% of initial Medicare spending.

Removing adverse selection causes surplus to increase considerably, but the results suggest that overall, imperfect competition contributes more to the observed distortions. Removing adverse selection eliminates between 36.0% and 44.4% of the difference in surplus between the observed and social planner's allocations. In markets with more intense competition, adverse selection plays a more important role in distorting outcomes. Where there is only one insurer, removing adverse selection eliminates, on average, between only 17.0% and 38.7% of the difference in surplus between the observed and efficient allocations. In markets with six or more insurers, removing adverse selection eliminates, on average, between 50.1% and 65.0% of the difference in surplus.

The remainder of the paper proceeds as follows. Section II defines adverse selection. Sections III and IV describe the M+C program and the data sources that I exploit. Sections V and VI describe the economic model and how it is estimated. Then Sections VII-IX discuss the empirical results and section X, the counterfactual simulations and welfare results. Section XI concludes.

II. Adverse Selection

Suppose we observe an insurance market that contains a set of insurers $j = 1, 2, .., J$ who each offer one insurance plan. These plans vary in their premiums (p_j) and the generosity of coverage offered (g_j)[5]. In this market, each consumer i has private information, $\theta_i \sim F(\theta)$, about his preferences, and using this private information, enrolls in the plan that maximizes expected utility:

$$EU_{ij} = \theta_i g_j - p_j + \varepsilon_{ij}$$

where ε_{ij} is a shock to the utility i receives from plan j. Note that heterogeneity in θ implies selection. Since $\dfrac{dEU_{ij}}{g_j} = \theta_i$, high θ consumers are more willing to trade income for benefits and will tend to enroll in generous insurance plans.

The presence of selection does not imply adverse selection. Selection is adverse only if consumers likely to enroll in generous insurance plans are also the costliest to insure.

Assume that an insurer offering coverage g to a consumer θ expects to incur costs $c(g,\theta)$. If, conditional on g, these costs are increasing in θ, then selection is adverse. Models of adverse selection often assume θ reflects consumers' unobserved health. But θ may reflect characteristics correlated with good health, such as risk aversion or cognitive ability[6]. Under these circumstances, insurers' expected costs are decreasing in θ and selection is advantageous to insurers. Therefore,

- selection is adverse to insurers if $\frac{dc(g,\theta)}{d\theta}|_g > 0$.
- selection is advantageous to insurers if $\frac{dc(g,\theta)}{d\theta}|_g < 0$.[7]

When selection is adverse ($\frac{dc(g,\theta)}{d\theta}|_g > 0$), market failures stem from two sources. First, if insurers are aware that selection is adverse, they have incentives to distort downwards the coverage they offer, to attract the least costly enrollees (Rothschild and Stiglitz, 1976). Second, prices in an unregulated market, regardless of the nature of competition do not effectively induce consumers to make efficient choices.[8] If i enrolls in plan j, the surplus generated from this choice, $s(i,j)$, is the sum of consumer and producer surplus:

$$\begin{aligned} s(i,j) &= EU_{ij} + p_j - c(g_j,\theta_i) \\ &= \theta_i g_j - c(g_j,\theta_i) + \varepsilon_{ij} \end{aligned}$$

Thus, i maximizes surplus if and only if $p_j = c(g_j,\theta_i)$ for all j. In the presence of adverse or advantageous selection this is not possible. For simplicity, assume that $\varepsilon_{ij} = 0$ for all ij so that for each j, there is an interval of consumers, $\left[\underline{\theta}_j^*, \overline{\theta}_j^*\right]$, that a social planner would allocate to j. This allocation can not be sustained in an unregulated perfectly competitive market. In such a market, premiums must reflect average costs, $p_j^* = \frac{\int_{\underline{\theta}_j^*}^{\overline{\theta}_j^*} c(g_j,\theta)dF(\theta)}{\int_{\underline{\theta}_j^*}^{\overline{\theta}_j^*} dF(\theta)}$. But if p_j is equal to p_j^*, $\underline{\theta}_j^*$ may not efficiently choose plan j since $p_j^* > c(g_j,\underline{\theta}_j^*)$. If $\underline{\theta}_j^*$ enrolls in plan j, he in effect, is forced to subsidize the relatively unhealthy consumers also enrolled in plan j. In response, $\underline{\theta}_j^*$ may inefficiently enroll in a plan with less generous coverage and a healthier pool of enrollees. Depending upon how severe the adverse selection is, competitive insurance markets can partially or fully unravel. Government intervention, in the form of subsidies for the generous plans can implement the efficient allocation (Feldman and Dowd (1982), Cutler and Reber (1998)).

When competition is imperfect, the allocation of consumers to insurance plans is still inherently inefficient because the premiums can not vary depending on enrollees' θ. However, the distortions to sorting may be larger or smaller than those that arise in a perfectly competitive setting. This will depend upon how plan markups vary with g.

A. Previous Empirical Literature

Given the theoretical importance of adverse selection, there exists an extensive empirical literature testing for its presence. Beginning with the positive correlation test introduced by Chiappori and Salanie (2000), much of this literature relates consumers' realized health care expenditures to their levels of insurance coverage. The logic of this test is as follows. If there is selection within an insurance market, then $\frac{dE[\theta|g_j]}{dg_j} > 0$. If plan j offers generous coverage, it expects to enroll high θ enrollees. Letting $c^R(g)$ denote realized costs, if selection is adverse, then $\frac{dc^R(g)}{dg} > 0$.[9] It is well known, however, that the presence of moral hazard can also induce this positive relationship. High θ consumers with generous coverage might be more costly to insure because they are inherently more costly (adverse selection), or because they face a lower price for health care (moral hazard).

In response, Finkelstein and Poterba (2006) suggest using "unused observables" to test for adverse selection. If there exists an observable characteristic x that is correlated with consumer demand and insurer costs, but not used when setting premiums, then there is adverse selection in x.[10][11] The "unused observables" test is robust to moral hazard but unable to detect the presence of market failures. Whether insurance markets operate efficiently depends upon the nature of selection along all consumer characteristics that determine demand for insurance. It is difficult to control for all of the characteristics that determine demand for insurance and when consumers have unobserved characteristics (for example, unobserved health), it is impossible.

B. Adverse Selection and Market Structure

The measure of adverse selection developed below overcomes the limits described above. I proceed in two steps. First, borrowing from the industrial organization literature on discrete choice, I specify and estimate a distribution of preferences similar to the one described above. This estimated distribution is intended to reflect all of the determinants of insurance demand, both observed and unobserved. Then, I specify a cost function $c(g, \theta)$ and use insurer behavior to directly estimate the partial derivative $\frac{dc(g,\theta)}{d\theta}|_g$, while controlling for unobserved heterogeneity across insurers and markets.

To identify and measure $\frac{dc(g,\theta)}{d\theta}|_g$, it is necessary isolate exogenous variation in the portion of $F(\theta)$ from which insurers offering a given g draw enrollees. To do so, I use variation in market structure. It is possible to detect adverse selection by examining the relationships between premiums (p) and benefits (g) across markets with different structures. The key idea underlying the strategy is that the ability of insurers to attract low cost enrollees depends on the benefits they offer, the benefits offered by competitors, and the number of competitors. To illustrate this strategy, consider the simplified markets presented in the three figures below.

Figures 1 and 2 represent two health insurance markets, the first with two and three insurers.

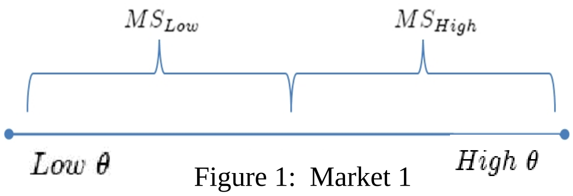

Figure 1: Market 1

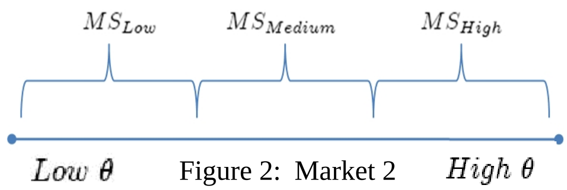

Figure 2: Market 2

Market 1 has plans with low and high generosity (g_{Low} and g_{High}). Market 2 has the same two insurers and third offering medium generosity (g_{Medium}). Assume that the presence of g_{Medium} is the only difference between markets 1 and 2. The horizontal lines represent the distribution of consumers' preferences and the arrows describe how consumers sort across plans. In market 1, g_{Low} and g_{High} each split the market, and g_{Low} attracts all of the low θ consumers whereas g_{High} attracts the high θ consumers. In market 2, the presence of a third insurers alters enrollees' sorting. g_{Low} and g_{High} are shown to have market shares close to $\frac{1}{3}$, drops that occur regardless of adverse selection. If adverse selection exists, however, the presence of of g_{Medium} differentially affects g_{Low} and g_{High}. If and only if adverse selection exists, compared to market 1, $g'_{Low}s$ enrollees become healthier and their costs go down. And compared to market 1, $g'_{High}s$ enrollees become less healthy and $g'_{High}s$ go up. Controlling for differences in the intensity of competition across the two markets, g_{Low} (g_{High}) will find it optimal to charge a lower (higher) premium in market 2 relative to market 1. The relationship between premiums and benefits offered becomes steeper in markets where the product space is more densely filled. This intuition is directly testable in the data and underlies my strategy to identify adverse selection.

Now, compare figures 1 and 3.

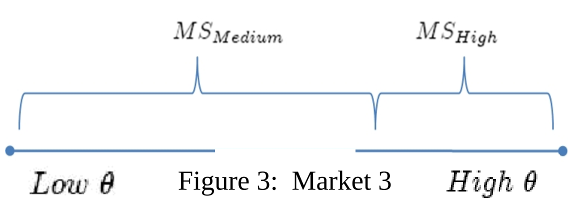

Figure 3: Market 3

As before, in both markets an insurer offers g_{High}. But whereas in market 1, g_{High} competes against g_{Low}, in market 3, g_{High} competes against g_{Medium}. The figures demonstrate that in market 3, $g'_{High}s$ enrollees have higher θs than in market 1. If there is adverse selection, this will translate into higher costs and controlling differences in demand and the nature of competition, a higher premium.

To summarize: Below I measure adverse selection by comparing the pricing behavior of insurance plans who offer the same insurance coverage, but operate in different markets and according to a sorting model, should attract different types of enrollees. The main challenge to this strategy is controlling for differences in costs across insurers and markets that can generate the same differences in pricing behavior. I discuss in great detail how I confront this challenge below. But to preview: First, in insurers' cost functions I include a wide range of observable variables that determine a market's structure and may be correlated with costs (for example, market size and average medical costs within a market). Second, I collect data on insurers' financial and operating characteristics which determine the types of coverage they offer but can reasonably be excluded from competitors' costs.

Though estimating a structural model with adverse selection requires strong assumptions about the economic primitives, the structure allows me to infer the impact adverse selection has on welfare. This paper, along with a few others[12] make up an emerging literature on the welfare consequences of adverse selection. Two features, however, distinguish this paper from those cited. First, I measure adverse selection using insurers' observed behavior, and not realized costs. A disadvantage to this approach is the required assumption that insurer behavior is optimal. But it allows for the analysis of larger insurance markets (such as privatized Medicare) where detailed cost data for all insurers in all markets is not readily available. Second, I make the more realistic assumption that insurance markets are imperfectly, rather than perfectly, competitive. Therefore, even in the absence of adverse selection, allocations will not be efficient. Below, I simulate what allocations will emerge if adverse selection is eliminated from Medicare, but the imperfect competition is retained. Comparing losses due to adverse selection and imperfect competition was not possible in the previous literature..

III. Medicare + Choice Program

Introduced in 1965, Medicare is the primary form of health insurance for the elderly and disabled, it is one of the federal government's largest programs and it constitutes a large portion of total health care spending.[13] In 1982, Congress passed the Tax Equity and Fiscal Responsibility Act, mandating the provision of managed care options to Medicare beneficiaries.[14] Since then, private insurers have played a continuous role in Medicare. I apply my framework to the years between 2000 and 2003 when the program governing privatized Medicare was called Medicare + Choice and HMOs were the dominant firm type.[15,16] Under M + C, Medicare beneficiaries could opt out of traditional Medicare and receive health insurance from a qualified private insurer. Insurers wishing to enroll Medicare beneficiaries signed contracts with the Center for Medicare and Medicaid Services (CMS) describing what coverage they would provide, and at what costs. A minimum set of benefits were required, essentially equal to the coverage included in traditional fee-for-service Medicare.[17] In exchange, the CMS made per-capita payments to each insurer under contract. Insurers had the option of providing additional benefits, such as (but not limited to) prescription drug, dental and vision coverage, as well as preventative care, in exchange for monthly premiums paid by enrollees.

Under M + C, it was believed that competitive pressures would induce insurers to submit proposed contracts that would achieve cost savings for the Medicare program. In addition, Medicare beneficiaries would benefit from coverage more generous than traditional Medicare.[18] This logic extended even to markets with few participating insurers, where insurers' plans still faced competition from government provided traditional Medicare.[19]

Policy makers are aware of potential selection problems within privatized Medicare and in particular, between traditional and privatized Medicare[20][21]. In response, the payments made to private insurers have increasingly been adjusted for enrollees' risk. Until 2000, government payments to insurers were set equal to 95% of the expected cost of treating a beneficiary within traditional Medicare, with adjustments made only according to enrollees' age, gender, and eligibility status. These adjustments, however, accounted for little of the variation in realized costs and did not reflect enrollees' health status.[22] The Balanced Budget Act of 1997 (BBA) called for expanded risk adjustment and dictated that payment rates to insurers reflect enrollees' lagged inpatient hospital experiences. This newer PIP-DCG risk adjustment model has been estimated to account for only 5-6% of variation in total health care costs (Medpac 2000). Within the sample period studied in this paper, 2000-2003, payments were weighted averages of the pre-BBA rate (90%) and the newer model (10%).

Contracts between the CMS and insurers were determined on a county and yearly basis after CMS payments rates were announced in each county. Once contracting between insurers and the CMS was completed, Medicare beneficiaries chose to remain in the traditional Medicare program or enroll in a private plan offered in their county. HMOs were not allowed to discriminate between beneficiaries within a market. Rather, insurers had to offer the same menu of plans to all individuals residing in the same county. Nor were they allowed to terminate a plan, increase premiums, or reduce coverage within a calendar year.[23]

Most beneficiaries that remained in traditional Medicare also enrolled in the voluntary part B component, and many enrolled in a supplementary Medigap plan.

IV. Data Sources and Descriptive Statistics

This paper combines data from multiple sources. Insurers' plan characteristics (and beneficiaries' choice sets) are retrieved from the Medicare Compare databases for the years 2000-2003. Market share data and individuals' plan choices are taken from the CMS State-County-Plan (SCP) files and the Medicare Current Beneficiary Survey (MCBS). HMO operating and financial characteristics are retrieved from the Weiss Ratings Guide to HMOs and Health Insurers. I extract market characteristics from the US Census and the American Hospital Association Directory. Variable definitions and construction of the final data set are discussed in the appendix.

The Medicare Compare Database is released each year to inform Medicare beneficiaries which private insurers are operating in their county, what plans they offer, and what benefits and costs are associated with each plan. For each plan, I collect information on dental coverage, vision coverage, brand and generic prescription drug coverage, and the copayments associated with prescription drugs, primary care doctor and specialist visits, and inpatient hospital admissions.

The SCP files provide market share data at the insurer-county-year level for all insurers in all counties. In each county-year, the SCP files also inform me how many residents are eligible for Medicare and the average CMS payment rate. Unfortunately, enrollment information is not provided at the plan level. For insurers offering multiple plans in a county, the market shares sum across the individual plan shares.

Tables 1 displays the distribution of insurers across markets and across time. The majority of markets have at least two insurers, but the overall trend in insurer participation is negative. HMO participation was at its strongest in 2000, when payment rate reductions put in place by the BBA began to take effect.[24] In 2003, the Medicare Modernization Act was passed, paving the way for expansions in privatized Medicare.

Table 2 displays the characteristics of markets that have different numbers of insurers. There are strong correlations between the number of insurers in a market and the government's payment rate, as well as between the number of insurers and market size.[25] Table 3 displays means and standard deviations of the plan characteristics I collect from the Medicare Compare database. It is through dental, vision, and prescription drug coverage that private insurers distinguish themselves from traditional Medicare.[26] The remaining variables proxy for cost-sharing requirements (i.e. copayments) in each plan.[27] In particular, there is extensive variation in premiums and the types of prescription drug coverage that are offered.

Table 4 displays a clear correlation between the number of insurers in a market and the amount of coverage typically offered. This reflects competition, payment rates, and perhaps costs. In markets with one insurer, the mean monthly premium is $68.64. In markets with six or more insurers, the mean premium is $14 per month. Similarly, in markets with more insurers, prescription drug, dental, and vision coverage become more common. Omitted from table 4 is information on inpatient hospital costs, which do not vary with the number of operating insurers.

The Medicare Current Beneficiary Survey tracks the behavior of a representative national sample of the Medicare population. Using the 2000-2003 versions of the MCBS, I match a sample of Medicare beneficiaries, many observed in multiple years, to specific insurance plans. The MCBS helps me overcome the lack of plan market share data in the SCP files.

MCBS administrative data informs me which M + C insurer each respondent is enrolled in. The Health Insurance section asks respondents specific questions about their Medicare managed care plan: whether dental, vision, or prescription drug coverage is provided, and what premiums are charged. Using answers to these questions, I identify which plan each respondent is enrolled in. The data set also informs me which county each respondent lives in. With the Medicare Compare Database, I construct choice sets for each MCBS respondent.

The MCBS also asks respondents detailed questions about their health status and lifestyles, and for those respondents remaining in traditional Medicare, gives a measure of total spending on health care. For example, the MCBS all respondents to rate their overall health on a scale from one to five, asks about experiences with illnesses such as diabetes, and respondents ability to perform activities such as bathing and walking one mile. Using responses to these questions and information on health care expenditures among enrollees in traditional Medicare, I am able to form a measure of health risk for those consumers opting into privatized Medicare.

Many types of insurers participate in M + C. Large, national companies such as Blue Cross and Blue Shield and Aetna offer plans in markets across the country. Smaller, regional HMOs offer plans in a few contiguous counties. To capture this heterogeneity, firm characteristics are extracted from the Weiss Ratings' Guides.

On a quarterly basis, the Weiss guides accumulate financial information from approximately 1500 United States HMOs and health insurers. This financial information includes total assets, capital, net income, etc. For all rated United States HMOs, Blue Cross and

Blue Shield plans, and large health insurers, the Weiss guide gives information on administrative expenses, enrollment levels, average medical expenses per enrollee, and physician network sizes.

Tables 5 and 6 describe the types of insurers participating in Medicare, and their characteristics. Most are Health Maintenance Organizations, but Preferred Provider Organizations and Private Fee for Service insurers are becoming increasingly common.[28]

Counties with distinct characteristics may impose different costs on insurers, and contain Medicare beneficiaries with distinct preferences. For each county, I retrieve population, income, and educational attainment data from the US Census. From the American Hospital Association Directory, I determine how many hospitals operate in each county. As noted above, there is a strong positive correlation between market size and the number of operating insurers. Larger markets attracting many insurers tend to be urban, which induces a negative correlation between entry and a market's average educational attainment, income, and proportion of citizens that are white.

V. Model

Insurers make decisions on a county and yearly basis and Medicare beneficiaries' choice sets include the plans available in the county they reside. On this basis, a market is defined as a county-year. The model below describes one market. Insurers decide whether to offer plans to Medicare beneficiaries, what coverage to provide, and what premiums to charge. Medicare beneficiaries choose between the plans offered and traditional Medicare.

A. Demand

Assume that consumer i, living in market m, obtains indirect utility u_{ijk} from plan k offered by insurer j and utility u_{i0} from traditional Medicare as follows:

$$\begin{aligned} u_{ijk} &= \theta_i g_{jk} - \alpha_m p_{jk} + \xi_{jm} + \varepsilon_{ijk} \\ u_{i0} &= \lambda_0 e_i + \chi_m + \varepsilon_{i0} \end{aligned} \tag{1}$$

where ρ, λ, α, ξ, χ, β and σ are the parameters to be estimated.

Each plan has observable characteristics x_{jk} and p_{jk}. p is the monthly premium. x describes $jk's$ coverage and cost sharing requirements. It includes doctor and hospital copayments and variables describing prescription drug, vision and dental coverage. The are collapsed into an index of overall generosity, $g_{jk} = x'_{jk} * \beta$.

Medicare beneficiaries have heterogeneous preferences for privatized Medicare and the benefits offered by these private plans. Preferences for g, $\theta_i = 1+\lambda_g e_i + v_i$ can be decomposed into an unobserved component, v_i which is drawn from the absolute value of a normal distribution with variance σ^2, and e_i, which proxies for observable health risk and is discussed further below. Individuals with large draws of ν gain more utility from plans with generous coverage and are more likely to enroll in a high g plan than individuals with low ν. Draws of ν are positive to ensure that the marginal utility of insurance coverage and disutility of premiums are positive.

The parameter λ captures the effect Medicare eligibles' observable health has on decisions to enroll in privatized or traditional Medicare and on preferences for generous private coverage conditional on enrolling in traditional Medicare. If consumers with poor observed health remain in traditional Medicare, the model will capture this with $\lambda_0 > 0$. To construct e, I use the MCBS which contains numerous measures of respondents' health. I regress these measures on total medical expenditures for services covered under parts A and B of traditional Medicare[29]. This regression, which yields an R^2 of .202 includes many of the same variables used for by the government when adjusting for payments to HMOs to reflect enrollee risk, and many not used (for example, questions about self-perceived health). These regression results, and a list of the variables used in forming e, are presented in table A1 in the appendix.

Finally, consumers' price sensitivity depends on income levels in market m. The amount of heterogeneity in preferences is determined by σ, heterogeneity in the characteristics that determine e, and income variation across markets.

For the outside good and each plan jk, i receives idiosyncratic type-1 extreme value shocks to utility, ε_{ijk}. Each insurer j also has unobserved characteristics, that provide ξ_j^munits of utility to all of $j's$ enrollees in market m. ξ can contain information the econometrician does not observe (for example, physician network quality or insurers' experiences in particular markets) and information the econometrician does observe (insurer type). A regression of ξ_j^m on the latter class of data can suggest which variables are correlated with high utility provision. Such a regression also yields estimates of χ_m. These fixed effects allow preferences for traditional Medicare to vary across markets.

B. Supply

Insurer j expects to incur a cost $c_{jk}\left(e, v, g, \Theta^{supply}\right)$ from providing coverage with generosity g to a Medicare beneficiary with preferences v and characteristics e. There is a fixed cost per enrollee, $FC_{jk}(e, v, \Theta^{Supply})$, that does not depend on g. The remaining variable costs of coverage, $MC_{jk}(e, v, \Theta^{Supply}) * g$, are linear in g :

$$c_{jk}(v, e, g, \Theta^{Supply}) = FC_{jk}(v, e, \Theta^{Supply}) + MC_{jk}(v, e, \Theta^{Supply}) * g \tag{2}$$
$$= \left[\pi_0 + \gamma_0 \nu + \mu_0 e + \delta_0' Z_{jm} + \psi_{jk}^0\right] + \left[\pi_1 + \gamma_1 \nu + \mu_1 e + \delta_1' Z_{jm} + \psi_{jk}^1\right] * g$$

where π, μ, γ, and δ are parameters to be estimated.

This cost function nests adverse and advantageous selection. If consumers with strong preferences for insurance coverage (i.e., those most likely to enroll in generous plans) are more costly to insure (adverse selection), γ_0 and γ_1 will be positive. If there is a negative relationship between preferences and costs (advantageous selection), γ_0 and γ_1 are negative. The signs and magnitudes of γ_0 and γ_1 play an important role in determining the effects of selection in equilibrium. I also include enrollees' observable health risk e in insurers' costs, though in principle, many of the observables that determine e are accounted for by Medicare's risk adjustment policies and therefore (since I only control for risk adjustment at the market level), should not enter my estimates of insurer costs.

Each insurer realizes two shocks to costs, ψ_{jk}^0 and ψ_{jk}^1 for each plan they offer. The shocks to FC and MC are perfectly observed by insurers, but not the econometrician, and are drawn from separate distributions that are *iid* and mean zero across insurers and markets.[30]

Finally, FC and MC depend on exogenous market and insurer characteristics, which are included in Z_{jm}. Z_{jm} controls for cost differences across markets and includes information such as average income, population, etc. In wealthier markets, for example, insurers may need to pay higher wages to local employees. Z_{jm} also includes the number of hospitals in m, yearly dummy variables, current CMS payment rates, as well as the pre-BBA payment rate. Prior to 1997, this rate perfectly reflected the government's costs of providing insurance to beneficiaries in traditional Medicare.

Z_{jm} includes insurers' financial and operating characteristics, such as lagged income and asset levels, administration costs, firm type, etc. Z controls for the possibility that heterogeneous insurers draw cost shocks from distributions with different means. For example, large insurers may be able to more effectively negotiate with pharmaceutical companies and therefore, have advantages offering generous drug coverage. Similarly, insurers with large physician networks may charge enrollees lower copayments for doctor office visits.

C. Discussion

My model allows consumers only one unobservable dimension of preferences for insurance coverage. If there are two consumers, i and i', with $v_i > v_{i'}$, then i has stronger preferences than i' for prescription drug and dental coverage. A more general environment might allow i to have stronger preferences for prescription drugs and i' to have stronger preferences for dental. I impose this restriction to simplify my analysis of insurer behavior. The restriction implies that the elements of x are perfect substitutes in each plan's market share function, since all consumers' willingness to trade between the elements of x depends on the same linear equation in β. In the appendix, I prove it is without additional loss of generality to assume insurers choose g, rather than x, for each plan.

A recent empirical literature emphasizes multiple dimensions of consumer heterogeneity, with adverse selection along some dimensions and advantageous selection along others. For example, risk aversion and cognitive ability, in addition to unobserved health, may determine preferences for insurance.[31] If consumers have T characteristics that determine their preferences (i.e. $v \equiv v(t_{i1}, t_{i2}, .., t_{iT})$), the elements of t may have different effects on insurers' costs. For example, t_{i1} might capture $i's$ unobserved health and t_{i2}, $i's$ risk aversion, which might be correlated with good health. In the appendix, I prove that under certain distributional assumptions on t and v, my model is equivalent to one with costs that depend on t (i.e. $c(g,t) = (\gamma_0' * t + \psi^0) + (\gamma_1' * t + \psi^1) * g)$. This result also follows from consumers having one dimension of unobserved preferences for insurance coverage.[32]

D. Assumptions

I describe in four stages the game that insurers and Medicare beneficiaries are assumed to play in market m. First, insurers decide how many plans to offer in each market, $n_j^m \geq 0$. Choosing $n_j^m = 0$ is equivalent to not entering. Next, after committing to n_j, the cost shocks ψ_{jk}^0 and ψ_{jk}^1 are realized for each plan k. Third, after observing competitors' entry decisions and ψ_{jk}^0 and ψ_{jk}^1 for each jk, each insurer simultaneously chooses g_{jk} and p_{jk} for each plan k. Insurer j chooses plan characteristics and premiums to maximize profits in market m. These choices satisfy a Nash equilibrium, which is assumed to exist. Medicare beneficiaries then observe these choices and sort across private plans and traditional Medicare to maximize expected utility. Letting $MS_{jk}(g, p, e, v \mid G, P, \Theta^{Demand})$ denote the percentage of Medicare beneficiaries with characteristics (e, v) who enroll in plan jk given the full set of plans offered, (G, P), realized profits are given by:

$$\int_e \int_v \sum_{k=1}^{n_j} MS_{jk}(g,p,e,v|G,P,\Theta^{Demand}) * [Subsidy_m + p_{jk} - c_{jk}(e,v,g,\Theta^{Supply})] dF(v|\sigma)dG(e)$$

Characterizing plan characteristics and premiums with a Nash equilibrium condition allows me to exploit insurers' first order conditions for g and p in the estimation routine. Similarly, the game's four stage structure implies a set of moment conditions useful in estimation. In particular, Stage 2 implies that conditional on Z, the cost shocks, ψ^0 and ψ^1, are independent of the number of plans and insurers in each market. It is necessary, therefore, that Z (which includes many insurer and market characteristics, including average medical costs within m) adequately control for differences in costs across insurers and markets.

VI. Estimation and Identification

This section provides a technical discussion of estimation and identification of the model described above. Let Θ_0 denote the true set of parameters. They are estimated using a simulated method of moments framework. Estimation is done in one step, but below, I discuss supply and demand separately. Details are left for the appendix.

A. Demand

Let Ω_m denote the set of private insurers offering plans in market m, and Ω_{mj} the set of plans offered by insurer j. The probability that an individual enrolls in plan jk in market m is given by:

$$Prob(jk|m,\Theta) = \int_e \int_v \frac{e^{u_{ijk}(v,e)+\xi_j^m}}{1 + \sum_{j' \in \Omega_m} \sum_{k' \in \Omega_{mj'}} e^{u_{ij'k'}(v,e)+\xi_{j'}^m}} dF(v|\sigma)dG(e) \tag{3}$$

In many discrete choice models, equation (3) and market share data are used to identify ξ. Remaining parameters are then identified using assumptions about the joint distributions of observed and unobserved product characteristics.[33] In my model, however, insurers observe ξ before choosing premiums and plan characteristics. Instead, my estimation strategy relies on two restrictions. First, there is no unobserved quality at the plan level, i.e., $\xi_{jk}^m = \xi_j^m$ $\forall$ jk and m. Second, individuals' vertical preferences, v, do not change over time. The first restriction allows me to identify θ, β, λ, and α. The second restriction allows me to identify σ. Below I discuss identification of ξ. Then I discuss identification of the remaining demand parameters.

The SCP files provide market share data at the insurer-market level. Let s_j^m equal insurer $j's$ market share in m. Following Berry (1994) and BLP (1995), for any θ, β, λ, α, and σ, there exists a unique ξ such that $s_j^m = \sum_{k\in\Omega_{mj}} Prob(jk|m,\Theta)\ \forall\ j\ and\ m.$[34]

If insurer j offers only one plan, then for any parameters, there is a ξ_j such that $j's$ predicted market share equals his observed market share. Thus, to identify β and α, I use the model to predict the choices made by MCBS respondents enrolled with insurers that offer more than one plan. Using Baye's rule, the probability that an individual enrolls in plan jk, conditional on enrolling with insurer j in market m, is given by $Prob(k|j,m,\Theta) = \frac{Prob(jk|m,\Theta)}{\sum_{k\in\Omega_{mj}} Prob(jk|m,\Theta)}$. Because there is no unobserved quality at the plan level (the first restriction), the model is forced to rely on the remaining demand parameters to predict within-insurer choices.

For an individual i enrolled with insurer j in market m, let d_{ijk}^m equal one if i enrolled in plan k and zero otherwise. d_{ijk}^m is a random variable with expected value $Prob(k|j,m,\Theta_0)$. Across jk, the only source of variation in d_{ijk}^m is from draws of ε_{ijk}. Since ε draws are independent of plan characteristics and premiums, at the true parameters the residual $\left[d_{ijk}^m - Prob(k|j,m,\Theta)\right]$ has expected value zero and is uncorrelated with plan characteristics and premiums. The moment condition below exploits these properties.

$$E[x_{jk} * (d_{ijk}^m - Prob(k|j,m,\Theta_0))|i,j,m] = 0 \tag{4}$$

Equation (4) states that at the true parameters, the model accurately predicts the mean level of x consumed by $j's$ population of enrollees. If the model predicts that too many of $j's$ enrollees purchase dental coverage, β is adjusted to correct this discrepancy. Similar moment conditions are constructed using other plan characteristics.[35]

To identify λ, I examine how sorting patterns across private plans and traditional Medicare depend on the distribution of observable health risk, $G(e)$. For example, I form moment conditions forcing the estimation routine to accurately predict the predicted expenditures of consumers enrolled in Medicare HMOs and traditional Medicare. Letting d_i^m indicate whether individual i, with health risk e_i, opted out of traditional Medicare, the population moment is given by:

$$E[e_i * (d_i^m - \sum_{jk\in\Omega_m} Prob(jk|m,\Theta_0))|i,m] = 0 \tag{5}$$

To identify λ_g, I form similar moment conditions that force the model to accurately predict how HMO enrollees with different predicted expenditures sort across insurance plans with different coverage levels.

To identify σ, I use the model to predict the choices made by MCBS respondents observed in multiple years. Let $Prob(jk \;\&\; j'k'|m, m', \Theta)$ denote the probability an individual is observed enrolling in plan jk in market m and $j'k'$ in market m' one year later. Because v does not change over time (the second restriction), $Prob(jk \;\&\; j'k'|m, m', \Theta) \neq Prob(jk|m, \Theta) * Prob(j'k'|m', \Theta)$. Let $d_{ijk,j'k'}^{m,m'}$ equal one if i is observed enrolling in plans jk and $j'k'$ in markets m and m', and zero otherwise. As before, $[d_{ijk,j'k'}^{m,m'} - Prob(jk \;\&\; j'k'|m, m', \Theta)]$ is a random variable whose only source of variation derives from draws of ε, and at the true parameters, has an expected value equal to zero. This motivates the following population moment:

$$(6) \qquad E[x_{jk} * x'_{j'k'} * (d_{ijk,j'k'}^{m,m'} - Prob(jk\&j'k'|m, m', \Theta_0))|i, m, m'] = 0$$

Condition (6) states that at the true parameter values, the model accurately predicts the covariance between characteristics of distinct plans chosen by the same individual in multiple years. As σ_g^2 increases, the model predicts more individuals with strong preferences for insurance coverage. These individuals are unlikely to enroll in a plan with little coverage in one year and a plan with generous coverage in the next. If, for example, MCBS respondents are rarely observed switching from plans with unlimited prescription drug coverage to ones with no drug coverage, this behavior can be explained by a large σ^2.

B. Supply

In each market the collection of premiums and plan characteristics, $\{g_{jk}, p_{jk}\}_{jk\in\Omega_m}$ are characterized by a Nash equilibrium. I assume insurers' first order conditions with respect to premiums and insurance coverage equal zero, $\frac{\partial Profit_j}{\partial g_{jk}} = \frac{\partial Profit_j}{\partial p_{jk}} = 0 \;\forall\; jk$. Thus, A firm offering n_j plans has $2 * n_j$ non-redundant, linear first order conditions in $2 * n_j$ cost shocks, ψ. For each choice of parameters, I invert the insurers' first order conditions to recover ψ_{jk}^0 and ψ_{jk}^1 for each jk. I construct moments from these recovered cost shocks to identify the parameters in costs, π, γ, μ, and δ.

If adverse (or advantageous) selection exists, insurers' marginal costs are directly affected by the preferences of their enrollees. The parameters of interest, γ and μ, measure this causal effect. Insurers' first order conditions and estimates of demand allow me to recover insurers' marginal costs, which are affected by adverse selection and insurers' unobserved cost shocks, ψ. Since insurers who realize different ψ have incentives to offer different types of plans, thereby attracting different types of consumers, I use exogenous variation in market structure to generate exogenous variation in costs related to adverse selection and identify γ and μ.

Consider a one-plan insurer's first order condition for generosity, g_{jk}:

$$(7)\quad p_{jk}\frac{dMS_{jk}(g,p,\Theta)}{dg_{jk}} = (\gamma_0+\gamma_1 g_{jk})\int_v v_g * \frac{dMS_{jk}(v|g,p,\Theta)}{dg_{jk}}dF(v|\sigma) + \gamma_1\int_v v_g * MS_{jk}(v|g,p,\Theta)dF(v|\sigma) + (\widetilde{\psi}_{0jk}+\widetilde{\psi}_{1jk} * g_{jk})\frac{dMS_{jk}(g,p,\Theta)}{dg_{jk}} + \widetilde{\psi}_{1jk}MS_{jk}(g,p,\Theta)$$

The left hand side of equation (7) is the marginal revenue to j from an incremental change in g_{jk}. The right hand side captures the marginal cost from an incremental change in g_{jk}. This marginal cost is the sum of three effects.[36]

First, increasing *generosity* changes the pool of enrollees that jk attracts. Increasing g attracts $\frac{dMS_{jk}(v|g,p,\Theta)}{dg_{jk}}dF(v|\sigma)$ new enrollees with preferences v. These new enrollees change the insurer's costs by $(\gamma_0+\gamma_1 g_{jk})\, v\frac{dMS_{jk}(v|g,p,\Theta)}{dg_{jk}}dF(v|\sigma)$. Integrating this indirect cost over v yields the first term in the right hand side of (7). If increasing g_{jk} attracts many *new* high v enrollees, and adverse selection is present, then costs will increase. If there is no advantageous or adverse selection ($\gamma_0=\gamma_1=0$), then changes in $jk's$ pool of enrollees do not affect costs.

Next, increasing g_{jk} has a direct effect on costs. It becomes more costly to insure each of $jk's$ enrollees (particularly high v consumers when adverse selection is present), even those enrolled in jk prior to the increase. This effect is captured by the second term in (7) and includes moral hazard. The third effect on marginal costs is determined by Z_{jm}, ψ^0_{jk} and ψ^1_{jk}.

The terms attached to γ_0 and γ_1 in (7), $(\int_v v * \frac{dMS_{jk}(v)}{dg_{jk}}$ and $\int_v v * MS_{jk}(v))$, are functions of elements in $\{g_{jk},p_{jk}\}_{jk\in\Omega_m}$ and therefore, correlated with ψ^0_{jk} and ψ^1_{jk} since insurers' plan characteristics are endogenous with respect to costs. To identify γ_0 and γ_1, I use the number of insurers and plans in each market (J and JK), ξ, and insurers' opponents' characteristics, Z_{-j} as instruments. The four stage game assumed in section (5) implies ξ, J and JK are independent of ψ_{jk}. I assume Z_{-j} is independent ψ_{jk}.[37]. The instruments are correlated with the endogenous terms because $\frac{dMS_{jk}(v)}{dg_{jk}}$ and $MS_{jk}(v)$ depend on the entire choice set, not just g_{jk} and p_{jk}. Consumer sorting depends on the number of plans and insurers in a market. And insurers' competitors' plan choices depend on Z_{-j}.

Intuition underlying these relationships was described in section 2. Markets with many insurance plans are less conducive to pooling amongst consumers with distinct v. As JK increases, the expected strength of $jk's$ enrollees' preferences, $\int_v v_g * MS_{jk}(v)$, increases more quickly in g_{jk}. Similarly, Z_{-j} affects $\int_v v_g * MS_{jk}(v)$ through g_{-jk}. If, for example, $Total\_Assets_{-j}$ is positively correlated with g_{-jk}, then $Total\_Assets_{-j}$ will be negatively correlated with $\int_v v{*}MS_{jk}(v)$.[38] An insurer whose competitor has high assets (and therefore, is likely to offer a generous plan), will tend to draw low v consumers. I exploit the population moments implied by the model, for example,

$$\begin{aligned} E[Mean(Z'_{-jm}) * JK_m * \psi^n_{jk}] &= 0 \quad \forall\, j,k,n \\ E[Mean(Z'_{-jm}) * \psi^n_{jk}] &= 0 \quad \forall\, j,k,n \end{aligned}$$

Given the assumptions in the model, the remaining parameters, π and δ, can be retrieved via an ordinary least squares regression. For any choice of γ and μ, let $\widetilde{\psi}^0_{jk}$ and $\widetilde{\psi}^1_{jk}$ be defined by $\widetilde{\psi}^n_{jk} = \pi_n + \delta'_n Z_{jm} + \psi^n_{jk}$ for $n = 0, 1$. After inverting insurers' first order conditions to recover $\widetilde{\psi}_{jk}$ (instead of ψ_{jk}), I regress $\widetilde{\psi}_{jk}$ on a constant and Z_{jm}.[39]

I estimate the model by constructing sample analogs of the population moments described above and using simulated generalized method of moments. A full list of moment conditions used, a discussion of simulation error and selection, as well as a summary of the estimation routine is available in the appendix.

VII. Estimation Results

A. Utility

The main specification for utility is given in equation (1). Table 7 displays estimates of the parameters entering this specification. The parameter estimates have the expected signs and are very precise. In particular, Medicare beneficiaries receive significant utility from prescription drug coverage. The parameters on *Generic PD* and *Brand PD Unlimited* are significant at the 0.1% level. The parameter attached to *Generic PD Unlimited* is significant at the 1% level.

Dental and vision coverage also provide positive utility. The payment variables are negative, as expected. The parameter attached to \$*Specialist* is not precisely estimated, but utility decreases by a significant amount when a plan's inpatient hospital and prescription drug cost-sharing increase.

Utility is decreasing in premiums, although less so in high income markets. To determine beneficiaries' sensitivity to premium increases, I calculate the mean semi-elasticity, $\frac{\partial MS_{jk}}{\partial P_{jk}} * \frac{1}{MS_{jk}}$. A \$1 increase in a monthly premium reduces enrollment, on average, by .64%.[40] I also calculate mean market share elasticities with respect to *generosity* and ξ ($\frac{\partial MS_{jk}}{\partial g_{jk}} * \frac{g_{jk}}{MS_{jk}}$

and $\frac{1}{n_j}\frac{\partial MS_j}{\partial \xi_j} * \frac{\xi_j}{MS_j}$). A one percentage increase in a plan's *generosity* and unobserved quality generate, on average, .76% and .29% increases in market share.

At the bottom of Table 7 are estimates of the parameters that determine the amount of heterogeneity in consumers' preferences. σ^2 is positive and estimated very precisely, suggesting there is heterogeneity in preferences for insurance coverage.

To interpret the economic importance of this heterogeneity, I calculate the willingness to pay for plan characteristics by consumers with different v. Table 8 provides estimates of the willingness to pay for each component of *generosity* at $50th$ percentile of e and the $10th$, $50th$, and $90th$ percentiles of the distribution of v. The median consumer is willing to pay up to $13.22 per month for dental coverage. An individual at the $90th$ percentile of the v distribution is willing to pay $7.73 more per month. Medicare beneficiaries are not sensitive to changes in doctor visit and inpatient hospital admission copayments. With respect to prescription drug insurance, a consumer at the $10th$ percentile is willing to pay $34.51 per month for prescription drug coverage that includes unlimited generic and branded drugs. An individual at the $90th$ percentile is willing to pay $81.83 per month for the most generous prescription drug package.

The inclusion of e in u_{ijk} and u_{i0} successfully controls for selection in observable health between traditional and privatized Medicare. Based on the distribution of observable characteristics and estimates of λ_g and λ_0, the predicted AB expenditures of enrollees opting into a private Medicare plan is $1833 per year. Among those remaining in traditional Medicare, predicted expenditures are on average $2057 per year, 12.2% higher

Table 4 indicates clear relationships between the number of insurers in a market, generous coverage, and lower premiums. Through improved coverage, lower premiums and plan variety, Medicare beneficiaries will extract more surplus from the M + C program in markets with more insurers. Table 9 provides estimates of consumer surplus per HMO enrollee and per Medicare beneficiary.[41] A representative Medicare beneficiary residing in a market with one insurer receives $19.85 in surplus per year from the M + C program, and in markets with more than six private insurers, $109.40 per year. Summing across all markets and years, the total contribution to consumer surplus is estimated to be $7.597 billion.[42]

B. Reduced Form Evidence of Adverse Selection

Using figures 1-3 in section 2, I argued that adverse selection can be detected by examining variation in insurers' costs (inferred from premiums) across markets with different structures. .The two specifications below test for the presence of adverse selection using this intuition.

$$(8) \quad p_{jk} = \sum_{n=2}^{8} \alpha_n 1[JK = n] g_{jk} + \sum_{n=1}^{5} \eta_n 1[J = n] + controls$$

$$(9) \quad p_{jk} = \gamma g_{jk} + \delta \sum_{n=1}^{JK} 1(g_n \leq g_{jk}) + \theta \sum_{n=1}^{JK} 1(g_n > g_{jk}) + \sum_{n=1}^{5} \pi_n 1[J = n] + controls$$

In (8), I regress plans' premiums on their coverage levels, g, allowing the coefficient on g to depend on JK, the number of plans available in a market. If adverse selection exists, we expect $\widehat{\alpha_n}$ to be increasing in n. In markets with more plans, there is less consumer pooling amongst consumers with different v. This hurts generous plans and helps less generous plans.

In (9), I regress premiums on coverage levels and variables that capture a plan's position within a market. $\sum_{n=1}^{JK} 1(g_n \leq g_{jk})$ is equal to the number of plans offering less coverage than jk. $\sum_{n=1}^{JK} 1(g_n \geq g_{jk})$ is equal to the number of plans offering more coverage. If adverse selection exists, conditional on g_{jk}, plans facing strong competition for low θ consumers (those with high $\sum_{n=1}^{JK} 1(g_n \leq g_{jk})$) will attract more costly pools of enrollees and be forced to charge a higher premium. On the other hand, plans with large $\sum_{n=1}^{JK} 1(g_n \geq g_{jk})$ face relatively weak competition for low θs, will attract less costly enrollees, and charge a lower premium.

Both specifications include controls for the number of insurers (J) operating in a market, capturing the overall effects of competition on premiums, as well as variables that control for cost differences across markets.

Table 10 presents estimates of (8) and (9). The results are consistent with adverse selection. The premium-generosity schedule does indeed become steeper in markets with more insurance plans. One-sided t tests confirm that α_2 is significantly less than α_3 at the 1% level. Similarly, α_3 is less than α_4, and α_5 is less than α_6 at the 5% levels. And conditional on a plan's coverage, the effect from an additional plan depends on the location of the additional plan. If a plan enters with less generous coverage, instead of more generous coverage, there is a larger positive effect on premiums. This difference in effects (δ versus θ) is highly significant.

The evidence in table 10 relies on comparisons between markets with different structures. Yet, market structure is endogenous: insurers choose which markets to enter and what types of plans to offer. The structural measures of adverse selection presented below attempt to be robust to this endogeneity. First, the cost and utility functions given in equations (1) and (2) included extensive controls for differences in demand and costs across markets which may affect entry and plan choices. Utility includes insurer-market fixed effects. Including such fixed effects in (2) would leave the model unidentified, but variables such as average health care costs, population, income, the number of hospitals, insurer financial and operating characteristics, and others attempt to control for differences in costs across markets and insurers The assumption that J and JK are independent of ψ is only made after conditioning on these variables.

Furthermore, my measure of adverse selection does not critically depend on the assumption that J and JK are independent of ψ. Rather, I detect adverse selection by looking for asymmetric effects of market structure (J and JK) on low and high coverage insurance plans. The relevant instruments are not J and JK, but rather, interactions between J, JK, and variables that determine whether a plan offers high or low amounts of g relative to its competitors. The use of these instruments, Z_{-j}, require the assumption that conditional on an insurer's own characteristics (his administrative expenses, his total assets, etc.), his competitors' characteristics do not directly enter his cost function.

C. Insurer Costs

The costs to plan jk from offering coverage g to a consumer with preferences for generosity, v, and observable health risk, e, is given in equation (2). Table 11 displays estimates of this specification.[43] The second row of Table 11 displays estimates of parameters relating to adverse selection ($\widehat{\gamma}_0$ and $\widehat{\gamma}_1$). Both parameter estimates, particularly $\widehat{\gamma}_0$, suggest the presence of adverse selection: beneficiaries with strong preferences for insurance are more costly to insure. A one unit increase in an beneficiary's type v causes a \$51.36 increase in insurers' fixed costs of providing care and a \$.55 increase in insurers' marginal costs of *generosity*. $\widehat{\gamma}_0$ is significant at the 1% level.

Across plans, the mean expected marginal cost to insurers from providing an additional unit of *generosity* to a consumer at the 10th percentile of $F(v|\widehat{\sigma})$ is \$30.70 per month. For a consumer at the 90th percentile of $F(v|\widehat{\sigma})$, this marginal cost is \$31.57 per month. These parameters imply that net of cost-sharing, dental coverage (unlimited prescription drug coverage) typically costs an insurer \$209.99 and \$215.94 (\$1130.62 and \$1162.66) per year for enrollees at the 10th and 90th percentiles of $Fv|\widehat{\sigma})$.

Adverse selection appears to operate primarily through fixed costs per enrollee, which include the costs of the mandatory coverage included in traditional Medicare but not the optional coverage included in g. For enrollees at the 10th and 90th percentiles of $F(v|\widehat{\sigma})$, mean fixed costs are \$6,249.42 and \$7,229.68 per year, which differ by 15.8%. It should be noted that these results, that a consumer's preferences for g have an economically larger impact on fixed costs, still imply the market failures associated with adverse selection. Insurers offering generous coverage are forced to charge high premiums that reflect their higher fixed costs of coverage.

The effect of beneficiaries' observable health risk (e) on insurers' costs is small and imprecisely estimated. This may reflect the presence of risk adjustment policies which are not controlled for. Many of the observable individual characteristics that enter e are used by the government to adjust insurers' payments for the risk of their enrollees. If a high e beneficiary is more costly to insure, but his insurer is compensated by the government for these extra costs and this extra compensation is not controlled for, than the estimation routine will determine that e does not impact costs.

The estimated coefficients attached to insurer and market controls in table 11 vary in statistical significance. Also, some conform with economic intuition and others do not. HMOs appear to have higher fixed and marginal costs than non-HMOs, and overall costs are lower for large insurers. Insurers employing more physicians per enrollee behave as if they pay higher costs per enrollee.

These controls' ability to explain costs, however, is unimportant relative to their ability to explain observed levels of insurance coverage and premiums. It must be true that the elements of Z_j are correlated with $j's$ plans' characteristics, conditional on market characteristics. If so, Z_{-j} can be used to generate the variation in q_{-j} and p_{-j} that identifies γ_0 and γ_1. Table A2, in the appendix, displays results from a regression of q_{jk} and p_{jk} on the market and insurer characteristics included in (2). Conditional on market characteristics, insurers' characteristics do help explain their premiums and chosen levels of insurance coverage.

VIII. Welfare Losses From Adverse Selection and Imperfect Competition

In this section, I use the estimated model and counterfactual simulations to learn about the effects of adverse selection and imperfect competition on welfare in Medicare. The use of both are necessary. The estimated model determines how a social planner would allocate beneficiaries across private plans and traditional Medicare, and the amount of surplus generated by any given allocation. Counterfactual simulations are necessary to distinguish between distortions due to adverse selection and imperfect competition. The estimated demand elasticities and observed relationships between premiums and entry suggest that insurers do possess market power. Therefore, if adverse selection is removed from the market, the social planner's allocation will not emerge; distortions due to imperfect competition will remain. To isolate the effects due to adverse selection, I simulate the equilibria that emerge after adverse selection is eliminated, but imperfect competition is retained.[44]

A. Definition of Efficiency

To find the surplus generated by placing individual i in plan jk or traditional Medicare, add profits to utility:

$$s_{ijk} = \theta_i g_{jk} - \alpha_m c_{jk}(e_i, v_i, g_{jk}) + \xi_{jm} + \varepsilon_{ijk} \tag{10}$$

$$s_{i0} = \lambda_0 e_i + \chi_m + \varepsilon_{i0} \tag{11}$$

Below, the efficient allocation is the one in which each beneficiary i is placed in the plan that maximizes s. Differences between market shares in the social planner's and observed equilibriums (and therefore, surplus) are driven only by the margins $p_{jk} - c_{jk}(e, v, g_{jk})$. If each consumer faces prices equal to their marginal costs, the social planner's allocation and observed market allocation coincide. When selection is adverse, this is impossible. Adversely selected markets are prone to unraveling because generous insurers charge premiums that reflect the costliness of their enrollee pool, and enrollees with distinct v face the same prices. Under the first simulation, generous insurers no longer charge premiums that reflect their unhealthy pool of enrollees and beneficiaries with distinct v face premiums reflecting the costs *they* impose on their insurers. The second simulation implements an extreme version of the risk adjustment that is currently in place. Insurers' costs reflect the coverage they offer but not their enrollees health, and beneficiaries do not face premiums that reflect their own health.

B. Counterfactual Simulations

Under simulation #1, the government compensates insurers according to the types of consumers they attract, each consumer pays a tax to fund this risk adjustment policy that depends on the costs they impose on their insurers, insurers maximize profits in response to this policy, and their chosen premiums satisfy a simultaneous move Nash equilibrium

If competition is perfect, the efficient allocation should emerge. Let v^*_{jk} and e^*_{jk} denote the expected values of plan $jk's$ enrollees' v and e that emerge in the social planner's allocation. Assume that the government implements perfect risk adjustment so that each insurer is compensated as if each of his enrollees were of type v^*_{jk} and e^*_{jk}. If plan jk in market m attracts a consumer i with characteristics v_i and e_i, the insurer receives a subsidy equal to:

$$S_m + c_{jk}(v_i, e_i, g_{jk}) - c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk}) \tag{12}$$

where S_m is equal to the initial subsidy.

This risk adjustment policy affects government expenditures and is funded by the

Medicare population. When consumer i enrolls in plan jk he alters government expenditures by $c_{jk}(v_i, e_i, g_{jk}) - c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})$ relative to the observed equilibrium If i is forced to internalize this change in expenditures through a tax or subsidy and competition is perfect, the efficient allocation emerges. Taking into account the taxes/subsidies, beneficiaries maximize utility:

$$(13) \qquad u_{ijk} = \theta_i g_{jk} - \alpha_m[p_{jk} + c_{jk}(v_i, e_i, g_{jk}) - c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})] + \xi_{jm} + \varepsilon_{ijk}$$

$$(14) \qquad u_{i0} = \lambda_0 e_i + \chi_m + \varepsilon_{i0}$$

Under this policy, insurer $j's$ costs for plan k are equal to $c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})$ for all enrollees, regardless of type. If competition is perfect, then $p_{jk} = c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})$ and the utility functions given by equations (13) and (14) coincide with the surplus functions (10) and (11). If the allocations that emerge under simulation #1 yield less surplus than the social planner's allocation, it is due only to imperfect competition.

The tax policy in simulation 1 may be prohibitively costly to implement and/or politically infeasible. Simulation #2 retains perfect risk adjustment but removes type-dependent taxation. Again, let v^*_{jk} and e^*_{jk} denote the expected values of plan $jk's$ enrollees' v and e that emerge in the efficient allocation. Furthermore, let $E[v_{jk}|G, P]$ and $E[e_{jk}|G, P]$ denote the expected type of consumer plan jk will attract given any distribution of plan characteristics and premiums in $jk's$ market. To induce insurer j to behave as if each plan k enrollee is type v^*_{jk} and e^*_{jk}, the government must inform each insurer j that they will receive a subsidy for each plan k equal to:

$$(15) \qquad S_m + c_{jk}(E[v_{jk}|G, P], E[e_{jk}|G, P], g_{jk}) - c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})$$

After this subsidy is implemented, insurer $j's$ costs for plan k are equal to $c_{jk}(v^*_{jk}, e^*_{jk}, g_{jk})$ regardless of enrollee type. A Nash equilibrium exists at a set of new prices such that each insurer is maximizing profits and the emerging allocations are consistent with initial expectations.

Before presenting the results of these simulations, two points must be emphasized. First, simulation #2 is not necessarily revenue neutral. The total surplus results below reflect any changes in government expenditures relative to the observed equilibrium. Second, under simulation #2, consumers with different (v, e) will not face premiums that reflect differences in their costs. Rather, when choosing between plans jk and $j'k'$, consumer i will face premiums that reflect differences between (v^*_{jk}, e^*_{jk}) and $(v^*_{j'k'}, e^*_{j'k'})$. Therefore, even if competition is perfect, the efficient allocation may not emerge.

C. Welfare Results[45]

First, I compare the observed allocation to the social planner's allocation. Observed surplus is substantially lower than what would be achieved by the social planner. Summing across all years, markets, and consumers, total surplus increase from \$27.92 billion to \$133.45 billion. This change is equivalent to 16.5% of the model's estimate of total Medicare spending by the government and insurers. Table 12 describes these changes by levels of entry. No discernible patterns appear. Across all entry levels, the social planner increases surplus per beneficiary by approximately \$1,000 per year. These results suggest that the presence of adverse selection and imperfect competition are causing privatized Medicare to provide far less surplus than is possible.

When I implement the counterfactual simulations, total surplus increases to \$74.76 billion under simulation 1 and to \$65.86 billion under simulation 2. These surplus increases are equivalent to 8.1% and 6.6% of total Medicare spending in the observed equilibrium. Table 13 below, provides total surplus per beneficiary under these two simulations by different levels of entry. At all levels of entry, allocations from simulation 1 than simulation 2.

Under simulation 1, 44.4% of the gap in surplus between the observed and social planner's allocation is eliminated. Under simulation 2, 36.0% of this gap is eliminated. This suggests that while adverse selection has a significant effect on welfare, imperfect competition plays an even larger role in explaining why the observed allocation generates less surplus than the social planner's allocation. Tables 14 examines the roles of adverse selection and imperfect competition in markets with different numbers of insurers. The table informs the reader, for each level of entry, what percentage of the gap in surplus between the observed and social planner's allocation is eliminated under a particular counterfactual simulation. Where these percentages are larger, adverse selection plays a more important role relative to imperfect competition in reducing welfare. For example, in monopoly markets, simulation 1 removes 38.7% of gap in surplus between the observed and social planner's allocations. Simulation 2 removes 17.0% of the gap. We can conclude, therefore, that in markets with only one insurer, imperfect competition plays a more important role in distorting allocations than adverse selection. As expected, in markets with greater numbers of insurers, and therefore, more intense competition, adverse selection plays an increasingly important role in distorting outcomes. In markets with four or more insurers, eliminating adverse selection eliminates at least 50% of the gap in surplus. However, even in markets with six or more insurers, where we might expect competition to be strong, allocations continue to be distorted by imperfect competition.

IX. Conclusion

This is the first paper to use observed insurer behavior to investigate welfare losses caused by adverse selection and imperfect competition in a health insurance market. I estimate a model of privatized Medicare in which insurers choose how much coverage to offer and what premiums to charge, and consumers vary in their preferences for insurance. The model allows consumers with different preferences to impose different costs on their insurers. To measure adverse selection, I use exogenous variation in market structure to identify a causal relationship between consumers' preferences and insurers' costs. This strategy allows me to infer whether consumers' preferences for insurance, which determine how much insurance they consume, contain information about their expected health. My measure is analogous to previous tests of adverse selection in the sense that I relate consumers' choices to insurers' costs. But unlike tests that rely only on consumer behavior, the measure of adverse selection developed here does not require knowledge of insurers' realized costs. And my model of consumer sorting and insurer plan choice allows me to quantify the welfare losses caused by adverse selection. Market failures caused by adverse selection begin with distortions in consumer and insurer behavior. Measurement of the resulting welfare losses, therefore, requires a model of consumer and insurer.

My empirical results confirm that consumers with strong preferences for generous insurance are more costly to insure. Using counterfactual simulations, I find that this adverse selection has substantial effects on welfare. I implement policies that combine perfect risk adjustment with preference dependent taxes and subsidies. The perfect risk adjustment policies, which are extreme versions of policies the government is currently experimenting with, adjust government payments to insurers according to the types of consumers they enroll, thereby eliminating the effects of adverse selection on insurers' costs and incentives. And after the taxes and subsidies are implemented, the premiums consumers pay reflect the costs they impose on their insurers. The simulated equilibria exhibit levels of surplus substantially higher than in the observed equilibrium. Between 2000 and 2003, total surplus associated with privatized Medicare increases by \$46.8 billion, or equivalently, by an amount equal to 8.1% of the total costs incurred by insurers and the government.

Adverse selection, however, is not the only source of market failures within Medicare. The privatized Medicare market is imperfectly competitive, and the welfare losses caused by insurers' market power are moderately larger than those due to adverse selection. The 8.1% increase in surplus accounts for only 44% of the difference in surplus between the observed and social planner's allocations. The remaining 56% is due to imperfect competition. In markets with six or more insurers, adverse selection accounts for over 60% of the market failures. In more commonly observed monopoly markets, however, adverse selection accounts for only 38.7% of the documented market failures. These results suggest that there are multiple avenues the government can pursue to improve the efficiency of privatized Medicare.

It should be pointed out however, that this analysis ignores the effects on welfare from plan manipulation by insurers (i.e., Rothschild & Stiglitz (1976)) in response to adverse selection. The framework developed does endogenize insurers' coverage choices and can, via simulations, quantify the amount of plan manipulation taking place. This is a likely avenue of future research.

Notes

*I am indebted to my advisors, Steven Berry and Philip Haile for their continual guidance, support, and encouragement. I also benefited from constructive comments made by numerous seminar participants and conversations with Fiona Scott-Morton, Marc Rysman, Jacob Gramlich, Katrina Jessoe, Ying Fan, and Alessandro Bonatti. Special thanks also go to Lyle Nelson at the Conressional Budget Office for greatly improving my understanding of Medicare.

†Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215; jlustig@bu.edu; http://people.bu.edu/jlustig/

[1]Medicare is a federal program that provides health insurance to the elderly and disabled.

[2]Fang, Keane, and Silverman (2004) and Finkelstein and McGarry (2006) find evidence of adverse selection in unobserved health, but advantageous selection in other characteristics, such as risk aversion, income, and cognitive ability. Advantageous selection exists when consumers enrolling in generous insurance have good unobserved health.

[3]Moral hazard takes place after contracting is completed. Thus, an insurer's costs that are attributable to moral hazard do not depend on market structure. They depend only on the insurer's plans' characteristics. Similar intuition was used by Cardon and Hendel (2001) to distinguish between adverse selection and moral hazard.

[4]If adverse selection is severe, insurers may be unable to simultaneously offer generous insurance and attract healthy (i.e. profitable) enrollees. Markets for generous insurance may unravel. Cutler and Reber (1998) provide evidence of such an unravelling in a Boston insurance market.

[5]Traditionally, in models of indemnity insurance, generosity of coverage might denote a coinsurance rate or deductible. For the managed care industry (where deductibles and coinsurance rates tend to always be low), it is more appropriate to think of generosity as capturing access to care and benefits such as prescription drug, dental, and vision coverage.

[6]For evidence of this, see Fang, Keane, and Silverman (2008) and Finkelstein and McGarry (2006).

[7]Note these these partial derivatives are with respect to θ and g fixed. This is particularly important if moral hazard exists. If moral hazard exists, consumers with more generous coverage consume more health care, or $\frac{dc(g,\theta)}{dg}|_\theta > 0$

[8]The model estimated in this paper is capable of quantifying both types of losses, but focuses on the second type.

[9]Cutler & Zeckhauser (1997) survey several papers documenting such positive correlations.

[10]If insurers are able to vary premiums with x, then the market failures associated with adverse selection will not occur.

[11]For example, suppose that insurers do not vary premiums with enrollees' education. If college graduates have higher demand for insurance and lower medical expenditures, then there is advantageous selection in education.

[12]Einav, FInkelstein, and Schrimpf (2007) measure the welfare losses due to adverse selection in an annuities market. Cullen, Einav, and Finkelstein (2008) and Budorf, Levin, and Mahoney (2008) estimate the welfare losses from adverse selection in markets for employer provided health insurance.

[13]Spending on Medicare in 2003 totaled \$315 billion and in 2002, Medicare accounted for 19% of total spending on personal health care and 2.6% of GDP (Medpac 2004).

[14]HMO participation in Medicare began in 1972. Participation, however, was minimal until TEFRA.

[15] Participation by Preferred Provider Organizations (PPOs) was beginning to develop and Private Fee for Service insurers (PFFS) were still in their infancy.

[16] The 2003 Medicare Modernization Act redefined and expanded the role played by private insurers, relabeling the program "Medicare Advantage" and adding a prescription drug component to the program.

[17] Traditional Medicare is comprised of two parts, A and B. Part A covers hospital services and enrollment is automatic. Part B covers outpatient services. Enrollment is not automatic and requires a monthly payment by enrollees, but it is heavily subsidized by general tax revenue. Virtually all Medicare beneficiaries enroll in Part B.

[18] "Over time, participating plans will be under competitive pressure to improve their benefits reduce their premiums and cost sharing, and improve their networks and services, in order to gain or retain market share" (Medicare Managed Care Manual, CMS)

[19] Most participating insurers did offer plans with coverage not provided by traditional Medicare. In markets with one insurer, for example, 61% of insurers offered some supplemental prescription drug coverage between 2000 and 2003.

[20] For a survey of the literature documenting selection bias between traditional and privatized Medicare see Hellinger & Wong (2000).

[21] While the focus of this study is selection within privatized Medicare, between HMO plans that offer varying levels of coverage, I do explicitly control for selection between traditional and privatized Mediare.

[22] Most studies estimated that this crude system of risk adjustment accounted for only 1-2% of the variation in realized costs (Medpac 2000, Pope, Kauttner, et al. 2004).

[23] They could, however, increase coverage or reduce premiums.

[24] The BBA reduced payments on average to insurers by 6% (Congressional Budget Office 1999).

[25] Maruyama (2006) finds that CMS subsidy rates are an important determinant of entry by Medicare HMOs.

[26] Town & Liu (2003) find that the main source of surplus gains from HMO participation in Medicare are utility gains from prescription drug insurance which was absent from traditional Medicare until 2006.

[27] Clearly, there are more plan characteristics in a typical insurance plan than listed. Unfortunately, reporting methods make comparison of many plans' characteristics difficult. These variables should serve as a reasonable proxy for each plan's amount of coverage.

[28] PPOs share characteristics with traditional Indemnity Plans and HMOs. HMOs tend to be restrictive about where enrollees can receive medical care and traditional Indemnity plans are not. Within PPOs, medical care can be received from all providers, but small penalties are imposed for out of network care. PFFS insurers provide coverage identical to traditional Medicare.

[29] This level of expenditures is constructed from H_PTARMB and H_PTBRMB in the Medicare files. This information, however, is only available for beneficiaries choosing not to opt out of traditional Medicare.

[30] The mean zero assumption is without loss of generality.

[31] See Fang, Keane, and Silverman (2006); Finkelstein and McGarry (2006)

[32] Since insurers can not discriminate against consumers with the same v, but distinct t, insurers' relevant cost functions reduce to $E[c_{jk}(t, g, \Theta^{Supply})|\, v_g(t)]$. This expected cost is equal to $c_{jk}(v_g(t_{i1}, t_{i2}, .., t_{in}), g, \Theta^{Supply})$ if $\frac{dE[t_m|v]}{dv} = \triangle_m$ for all v and m.

[33] See Berry, Levinsohn, and Pakes (2003), who assume that observed product characteristics are mean independent of ξ. Other studies of the Medicare HMO market, such as Town & Liu (2003) and Maruyama (2006) make similar assumptions.

[34] ξ is unique up to a normalization. I normalize the mean utility of traditional Medicare to be zero.

[35] The integrals used to construct $Prob(k|j,\Theta)$ do not have closed formed solutions and must be simulated using a sample of draws from $F(v|\sigma)$. The induced simulation error enters these moments linearly, and therefore does not generate any bias. The sample moments' variances are affected, however. Corrections to these variances are discussed in the appendix.

[36] For clarity, in (7), $\widetilde{\psi}^z_{jk} = \pi_z + \delta'_z X_m + \lambda'_z Z_j + \psi^z_{jk}$ for $z = 0, 1$. And, I have temporarily omitted observable health e from the model.

[37] ψ^0_{jk} and ψ^1_{jk} are realized conditional on Z_j and X_m, and therefore uncorrelated with Z_j and X_m. Therefore, it seems reasonable to assume Z_{-j} is also uncorrelated with ψ^0_{jk} and ψ^1_{jk}.

[38] Results from reduced form regressions of Z_j on g_{jk} and p_{jk} are provided in table A2 .

[39] This method is not asymptotically efficient. A weighted IV regression using X_m, Z_j, and Z_{-j} is implied by the GMM criterion function. In practice, this regression has little effect on parameter estimates and the standard errors.

[40] Within the same market, Town & Liu (2003) estimate a mean semi-elasticity with respect to insurers equal .9%. This elasticity is comparable to mine since the average insurer offers 1.64 plans. Buchmueller (2000) find a premium semi-elasticity of -.7%. Atherly, Dowd, and Feldman (2003) find -.7% and -.4%. Both of these studies also examine Medicare beneficiaries.

[41] Relative to a market with only the outside good, the expected consumer surplus a representative Medicare beneficiary receives in a market m from the presence of M + C insurers is $\int \frac{1}{(1+\theta_p I_m + v_p)*\alpha} log\left(1 + \sum_{jk\in\Omega_m} e^{u_{jk}(v)}\right)$.

[42] Town & Liu (2003) estimate national consumer surplus from Medicare HMOs to $4.061 billion in 2000. Maruyama obtains similar estimates of $3.881 and $3.963 billion in 2003 and 2004.

[43] At the estimated parameters, I evaluate insurers' second order conditions for g and p to determine whether insurers are in fact behaving optimally. All of insurers' choices maximize profits locally, and in most cases, over a wide range of feasible (g, p) combinations.

[44] In related papers, Bundorf, Levin, and Mahoney (2008) and Einav, Finkelstein, and Cullen (2008), the authors assume that the observed equilibrium does not suffer from distortions due to imperfect competition. Given this, the effects of adverse selection on welfare can be calculated as differences in surplus between the social planner's and observed allocations.

[45] These result represent simulated equilibria for all markets. The simulated Nash equilibria were found by iterating on insurers' best responses. Start with the observed premiums and allow each insurer to simultaneously best respond given the new policy. Update premiums and allow insurers to simultaneously best respond again. Continue this process until convergence. Convergence is not guaranteed, but occurred within each market quickly.

REFERENCES

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

Atherly, A., Dowd, B., Feldman, R. "The Effects of Benefits, Premiums, and Health Risk on Health Plan Choice in the Medicare Program." Mimeo, Division of Health Services Research and Policy, University of Minnesota, 2003.

Berry, S., "Estimating Discrete-Choice Models of Product Differentiation." RAND Journal of Economics, Vol. 25 (1994), pp. 242-262.

Berry, S., Levinsohn, J., and Pakes, A. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Vehicle Market," Journal of Political Economy, 112(1), 68-104.

Bundorf, K., Levin, J., and Mahoney, N. "Pricing and Welfare in Health Plan Choice." Working Paper, 2008.

Cardon, J. and Hendel, I. (2001) "Asymmetric Information in Health Care and Health insurance Markets: Evidence from the National Medical Expenditure Survey." Rand Journal of Economics, 32(3), 408-427.

Chiappori, P., and Salanie, B. (2000), "Testing for asymmetric information in Insurance Markets,' Journal of Political Economy 108(1), 56-78.

Cullen, M. Einav, L. and Finkelstein, A. "Estimating Welfare In Insurance Markets Using Variation in Prices.", Working Paper, 2008.

Einav, L., Finkelstein, A., and Schrimpf, P. "The Welfare Cost of Asymmetric Information: Evidence from the UK Annuity Markets.", Working Paper, 2008.

Cutler, D.M. and Reber, S.J. 1998 "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection." Quarterly Journal of Economics, Vol. 113 (1998), pp. 433-466.

Cutler, D.M. and Zeckhauser, R.J. 1997 "Adverse Selection in Health Insurance." NBER Working Paper 6107.

Fang, H., Keane, M., and Silverman, D. (2006), "Sources of advantageous Selection: Evidence from the Medigap insurance market," mimeo, Yale University.

Finkelstein, A., and McGarry, K. (2006) "Multiple dimensions of private information: Evidence from the long-term Care Insurance Market," American Economic Review 96(4), 938-958.

Finkelstein, A. and Poterba, J. (2006), "Testing for Adverse Selection with 'Unused Observables'." NBER Working Paper No. 12112

Hellinger, F. and Wong, H. (2000) "'Selection Bias in HMOs: A Review of the Evidence." Medical Care Research and Review 57(4) 405-438.

Maruyama, S., (2006) "Welfare analysis incorporating a Structural Entry-Exit model: A Case Study of Medicare HMOs", Working Paper, Northwestern University.

McFadden, D., and Newey, W. (1994) "Large Sample Estimation and Hypothesis Testing", Handbook of Econometrics. Volume 4, Ch. 36.

Medicare Payment Advisory Commission, (2000-2004). "Report to the Congress: Improving Risk Adjustment in Medicare." Washington, DC: MedPac.

Rothschild, M., and J.E. Stiglitz. (1976) "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect information.", QJE, 90(1), 630-649.

Town, R. and Liu, S., (2000) "The welfare impact of Medicare HMOs," Rand journal of Economics, 34(4), 719-736.

Table 1: Total Number of Markets

# Insurers	2000	2001	2002	2003	Total
1	347	347	430	386	1510
2	152	167	132	141	592
3	114	77	60	76	327
4	62	22	16	23	123
5	31	13	4	5	53
6	17	5	2	4	28
>6	12	9	7	8	36
Total	735	640	651	643	

Table 2: Market Characteristics

# Insurers	Payment Rate	Outside Mkt. Sh.	# Plans in Mkt	Mkt Population
1	513.70	92.3%	2.06	121.92
2	526.78	85.5%	3.79	294.40
3	545.49	77.2%	6.96	442.00
4	537.97	73.5%	6.36	567.10
5	582.84	71.4%	8.95	663.70
6	603.54	72.8%	11.71	1395.16
>6	628.69	65.4%	12.70	3128.84

Notes: Payment rate is equal to the base subsidy paid per month/enrollee to HMOs by the CMS. Outside Mkt Share equal to the percentage of enrollees remaining in traditional Medicare. # Plans is equal to the total number of plans offered by insurers. Population is in 1000s.

Table 3: Plan Characteristics

	Mean	Std. Dev
Premium	$56.72	49.68
$ Specialist	$13.12	9.44
$ Inpatient Hospital Admission	$90.55	214.53
Vision	83.10%	3.70
Dental	19.78%	3.98
Generic PD	61.32%	4.87
Brand PD	50.05%	5.00
Generic PD Unlimited Coverage	34.17%	4.74
Brand PD Unlimited Coverage	6.44%	2.45
$ Copay Brand PD	$10.56	15.01

Notes: Premium is paid monthly. $Specialist and $Inpatient Hospital are paid when an enrollee visits a specialist or is admitted to an inpatient hospital. Vision, Dental, and Prescription Drug (PD) are dummy variables, indicating whether any coverage is offered. Unlimited coverage variables indicate if there are limits to prescription drug coverage. $ Copay Brand PD is enrollee's cost for a 30 day supply of a brand prescription drug. Payments all in 2000 dollars

Table 4: Mean Plan Characteristics by # Insurers

# Insurers	Premium	Dental	Vision	$ Specialist	Generic PD	Brand PD
1	$68.64	11.21%	71.98%	$12.88	46.48%	36.59%
2	$57.55	20.23%	83.23%	$13.62	57.00%	46.02%
3	$62.18	20.96%	88.34%	$14.26	67.21%	53.53%
4	$42.91	22.47%	90.73%	$12.25	74.86%	62.64%
5	$45.19	24.19%	94.19%	$12.30	82.79%	73.26%
6	$32.56	30.03%	91.47%	$11.77	83.28%	74.74%
>6	$14.03	45.45%	94.17%	$10.82	80.42%	68.07%

Table 5: Insurer Types

	HMO		PFFS		PPO	
Year	%	#	%	#	%	#
2000	92.85%	2200	0.00%	0	7.15%	134
2001	93.62%	1273	1.93%	16	4.45%	63
2002	87.17%	675	1.83%	11	11.00%	116
2003	80.52%	868	0.96%	7	18.52%	259

Notes: Table gives distribution of firm types in those markets used in estimation.

Table 6: Mean Insurer Characteristics

	Mean	Std. Dev
Total Assets	$374.63	791.78
Net Income	$10.82	42.12
# Member Physicians / Enrollee	0.473	5.20
% Administrative Expenses	9.88%	0.94
Medical Expenditures / Enrollee	$140.49	189.25
% Business Medicare	22.38%	2.17

Notes: All data reported at the insurer level, not the insurer-market level. Assets and Net Income reported in $Millions. #Physicians/Enrollee is equal to the total number of member physicians divided by the total number of enrollees. Administrative expenses are given as a percentage of total premium income. Medical Expenditure/Enrollee the average amount spent on each enrollee per month. %Business Medicare is percent of enrollees who are Medicare beneficiaries. All variables lagged one year.

Table 7: Utility Paramater Estimates

	(1)
Plan Characteristics	
Dental	0.5677 (.1488)***
Vision	2.6391 (.5129)***
$ Specialist	-0.007 (.2815)
$ Inpatient Hosp Admission	-0.9767 (.3087)***
Generic PD	1.1618 (.259)***
Brand PD	0.0916 (.3052)
$ Brand PD Copay	-0.0001 (.0087)
Generic PD Unlimited	0.2934 (.1628)*
Brand PD Unlimited	0.698 (.3197)*
Premium	-3.7694 (.8632)***
Vertical Preferences	
$Sigma^2$ Generosity	1.101 (.3138)**
Predicted Expenditures	0.0760 (.849)
Market $Income_{Premium}$	-0.0008 (.0003)***
Outside Good	
Predicted Expenditures	(.3229) (.1471)*

Notes: Standard errors in parentheses. * Significant at 5% level, ** at 1% level, *** at 0.1% level

Table 8: Willingness to Pay Estimates

	10th Percentile	Median	90th Percentile
Dental	$8.83 (2.32)	$13.22 (3.18)	$20.95 (4.82)
Vision	$40.58 (10.67)	$60.53 (13.49)	$95.63 (18.86)
$ Specialist	$0.02 (.24)	$0.02 (.35)	$0.03 (.55)
$ Inpatient Hospital Admission	$0.03 (.01)	$0.05 (.01)	$0.07 (.02)
PD Generic	$17.45 (3.52)	$26.16 (4.67)	$41.48 (7.02)
PD Brand	$2.01 (4.89)	$3.01 (7.30)	$4.75 (11.56)
$ Copay Brand	$0.02 (.13)	$0.03 (.20)	$0.05 (.32)
PD Generic Unlimited	$4.53 (3.09)	$6.76 (4.49)	$10.67 (6.99)
PD Brand Unlimited	$10.52 (5.26)	$15.74 (7.65)	$24.93 (11.97)

Notes: WTP estimates give Monthly $ value to consumer from $1 reductions in payment variables and $value to consumer from full coverage of indicator variables. It is assumed for each percentile, consumers have median preferences for premiums. Standard errors in parentheses. They are constructed after taking parameter draws from their estimated distributions.

Table 9: Consumer Surplus Estimates

# Insurers	CS/Medicare Beneficiary	CS/M+C Enrollee
1	$20.02	$239.13
	(10.69)	(76.86)
2	$51.33	$291.83
	(36.16)	(65.46)
3	$81.18	$322.87
	(44.8)	(66.88)
4	$105.10	$358.49
	(50.45)	(71.70)
5	$109.74	$374.09
	(52.18)	(71.88)
6	$113.49	$369.41
	(52.75)	(72.16)
>6	$156.95	$427.28
	(55.43)	(75.74)
Overall	$89.99	$339.08
	(50.75)	(73.64)

Notes: Column 1 gives yearly value in year 2000 dollars to a representative consumer from living in market with N insurers instead of zero insurers. Column 2 gives the total surplus per consumer actually enrolling in a private insurer. Values are population weighted means. Standard errors of means, calculated by taking draws from estimated distribution of parameters, are in parentheses.

Table 10: Reduced Form Evidence of Adverse Selection

	Equation 9	Equation 10
G*(JK=2)	-1.251	
	(0.364)**	
G*(JK=3)	0.56887	
	(0.20633)**	
G*(JK=4)	0.902	
	(0.151)**	
G*(JK=5)	1.052	
	(0.128)**	
G*(JK=6)	1.282	
	(0.098)**	
G*(JK=7)	1.091	
	(0.100)**	
G*(JK>7)	1.237	
	(0.063)**	
J=1	93.916	86.19
	(6.228)**	(6.293)**
J=2	80.374	77.792
	(4.891)**	(4.991)**
J=3	75.013	78.142
	(5.066)**	(5.209)**
J=4	54.687	60.577
	(5.267)**	(5.347)**
J>4	36.827	43.313
	(5.681)**	(5.909)**
G		3.54
		(0.512)**
#JK with less G		3.3319
		(0.27490)**
#JK with more G		1.272
		(0.251)**
Observations	5618	5618
R-squared	0.616	0.603

Notes: Dependent variable in both regressions is a plan's monthly premium. "G" is equal to plan's overall level of generosity, derived from estimates of demand parameters. "JK" is equal to the number of plans in a market. "J" is equal to the number of insurers. Control variables include insurer status, year dummies, and average Medicare costs within a market. Standard errors in parentheses. * indicates significance at the 5% level. ** indicates significance at the 1% level.

Table 11: Cost Parameter Estimates

	Fixed Cost	MC of Generosity
Constant	8.641	10.078
	(7.62)	(4.27)**
v_g	51.3567	0.5437
	(17.563)**	(2.23)
Predicted Costs (e)	0.5021	0.0911
	(171.8)	(55.8)
AB Rate / 10	9.288	9.535
	(.1391)***	(3.72)**
AB Rate 1997 / 10	-0.514	-0.072
	(.1097)***	(.083)
Population / 1000	0.014	-0.031
	(.017)	(.065)
Per Capita Income / 1000	10.209	-0.0001
	(1.209)***	(.01)
# Hospitals	-0.3779	-0.2801
	(.1329)**	(.072)***
2001 Dummy	9.777	-1.757
	(1.667)***	(.9907)*
2002 Dummy	15.353	-1.479
	(2.033)***	(1.208)
2003 Dummy	24.329	-0.1608
	(1.779)***	(1.057)
HMO Dummy	-1.6392	0.0227
	(2.658)	(1.57)
Total Assets ($million)	-3.3942	0.0262
	(.947)**	(.56)
Net Income ($million)	0.0285	-0.001
	(.0186)	(.011)
# Network Physicians / Enrollee	1.1263	0.057
	(.1537)***	(.0914)
% Administrative Expenses	-13.13	-0.023
	(1.087)***	(.6457)
Medical Expenditures / Enrollee	-0.004	0.0001
	(.006)	(.003)
% Business Medicare	-1.712	0.165
	(.3759)***	(.223)

Notes: Insurer controls lagged one year. AB Rate and AB Rate 1997 are current and 1997 government payment rates to insurers. Other variables defined in tables 2 and 4. *,**,*** indicate significance at 5%,1%, and 0.1% levels. Standard Errors are in parentheses.

Table 12: Total Surplus / Medicare Beneficiary

# Insurers	Observed Equilibrium	Social Planner's Allocation	Change ($)	Change (% of Costs)
1	$96.15	$1,039.86	$943.71	14.8%
2	$219.33	$1,370.86	$1,151.53	18.4%
3	$311.84	$1,405.41	$1,093.57	17.1%
4	$390.03	$1,400.80	$1,010.77	15.8%
5	$411.84	$1,582.18	$1,170.34	17.1%
6	$418.20	$1,496.45	$1,078.25	16.3%
>6	$539.18	$1,442.97	$903.79	13.0%

Note: Calculations are per Medicare beneficiary and per year. Columns 1-3 are in 2000 dollars. Column 4 gives changes in surplus as a percentage of market share weighted costs. Means are weighted by number of beneficiaries in each market.

Table 14: Adverse Selection v. Imperfect Competition

# Insurers	Simulation #1	Simulation #2
1	38.7%	17.0%
2	44.0%	34.7%
3	49.8%	41.5%
4	53.3%	44.2%
5	56.1%	45.9%
6	60.0%	49.5%
>6	65.0%	50.1%

Note: Values in table give the percentage of the gap in surplus across the observed and social planners allocation that is eliminated after the counterfactual simulations are implemented. Calculations are done by levels of entry. For each level of entry, calculation is a population weighted mean across markets.

Table 13: Simulated Total Surplus / Beneficiary

# Insurers	Simulation #1	Simulation #2
1	$371.14	$291.00
2	$693.83	$616.88
3	$821.97	$747.94
4	$935.84	$842.09
5	$1,027.39	$931.25
6	$1,048.51	$948.91
>6	$1,131.11	$993.19

Note: Total Surplus amounts are per year and per beneficiary under allocations that emerge under simulations described in the text. Results are in year 2000 millions of $. Standard deviations of means are in parentheses.B84

NOT FOR PUBLICATION

X. Appendix

A. Data

The main sources of data are acquired from the Center for Medicare and Medicaid Services. The data primarily comes from three sources: the yearly Medicare Compare database, the December version of the CMS's State-County-Plan files, and the MCBS survey. In the CMS data files, there is a clear distinction between organizations (insurers), contracts, and products (plans). The organization is the broadest of the three and is the company responsible for overseeing each contract. For example, Pacificare, Aetna, and Humana are all organizations often observed within the data. In the Medicare program, each organization enters into one or more contracts with the CMS for the purpose of providing health care to Medicare beneficiaries. Each contract is assigned a unique contract number by the government. Often times, organizations enter into multiple contracts with the government, with different contracts serving different markets. Other times, the same contract extends across markets. Within each contract, organizations offer, and Medicare beneficiaries can enroll in, multiple products. Products within a given contract can vary along multiple dimensions, including premiums, copayments, coverage types, etc. In the main text, I assume organizations' behavior is independent across markets. I refer to a contract as a firm or insurer and a product as a plan.

Identifying Market Entrants.—To identify insurers, I focus on the contract level. The State-County-Plan files disaggregate enrollment data at the contract level into enrollment data at the county level. For each Medicare HMO contract, I observe which counties have a positive number of contract enrollees. If this positive number is greater than ten, I observe actual enrollment information. Using the SCP data alone, however, greatly overstates entry by Medicare HMOs. When initially enrolling in a Medicare HMO, beneficiaries are only able to choose among those plans offered within their county. But once enrolled, beneficiaries can remain in a particular plan even after changing residences. Thus, within the SCP files, there are many counties that contain information on the enrollment by Medicare beneficiaries in a particular contract, despite that contract not being available within the county. To alleviate this problem, I eliminate all insurers with less than a 1% county market share. In addition, to be included as a participant in a county, an insurer must be listed in the Medicare Compare Database. To enter the final sample, a contract must be listed as being available in the Medicare Compare database, and have a market share of at least one percent.

Omitted Plan Types.—Under the Medicare + Choice program, several types of private health insurance plans other than Medicare HMOs are eligible to contract with the CMS to provide coverage to Medicare beneficiaries. Other types of organizations include Preferred Provider Organizations (PPO), Provider Sponsored Organizations (PSO), Private Fee For Service Plans (PFFS), and cost contract HMOs. From a beneficiary's point of view, these plans are full-fledged alternatives to Medicare HMOs as long as they are available in the beneficiary's market. They are included in the model.

Other types of plans, however, are excluded from supply and demand. These types include Health Care Prepayment Plans (HCPP), which cover only outpatient services, Programs of All-Inclusive Care for the Elderly (PACE), which are combination programs with Medicaid that provide comprehensive community and medical services (to be enrolled in a PACE plan a Medicare beneficiary needs to be certified as eligible for nursing home care by the appropriate state agency), and demonstration plans (DEMO) which are designed to evaluate the effects and impacts of various health care initiatives.

Plan Characteristics.—Definitions of variables entering Medicare beneficiaries utility are provided below. All product characteristics are extracted from the Medicare Compare database.

$Premium$: Premium information is provided at the monthly level. They do not include the required payment for Medicare Part B, which is charged to all Medicare HMO enrollees. For a few plans, premium information was unavailable. These plans were dropped from the final sample.

$\$InpatientHospital$: Required payments for inpatient hospital stays varied across plans in structure. Some described required copayments for days $0-20$ and $20-100$, others had additional tiers. This variable is equal to the charged copayment upon admission to an inpatient hospital. For some plans, this information is explicitly provided in the Medicare Compare Database. For others, the day 1 copayment is used.

$\$Specialist$: Required copayment for each visit to a specialist. Across plans, the format of this data is constant.

$Dental$: This is a dummy variable indicating whether any supplementary dental coverage is described in the Medicare compare database. Traditional Medicare offers no dental coverage. For some plans, additional information on dental coverage was provided, such as annual deductibles, dentist copayments, etc. But much of this information was incomplete and its structure varied extensively across plans. If no mention of Dental coverage was made, I assumed there was none.

$Vision$: This is a dummy variable indicating whether any supplementary vision coverage is described in the Medicare compare database. Traditional Medicare offers very little coverage of vision services. The formats of the data provided was very similar in nature to *Dental*.

PrescriptionDrugs: Five prescription drug variables were used in demand. First, *PDBrand* and *PDGeneric* are dummy variables indicating whether each plan offered any form of generic and branded prescription drug coverage. Between 2000 and 2003, traditional Medicare offered no such coverage. $\$BrandPD$ is equal to the copayment for a 31 day supply of a brand prescription drug. Some plans described this copayment using different lengths of time. For these, all copayments were prorated to 31 days. The generosity of provided prescription drug coverage is captured by $UL\_BrandPD$ and $UL\_GenericPD$. Both are dummy variables indicating whether generic and brand coverage are unlimited. Most plans offered only limited drug coverage, for example $1000 annually. Because of variation in the format of coverage data across plans, it was difficult to construct additional variables detailing the generosity of any provided drug coverage. If for each plan, any mention is made of unlimited generic or brand drug coverage these indicator variables are set to one.

B. Insurers' Choice Variables

The model assumes insurers choose generosity levels, g, and not the determinants of generosity, x. Note that in insurers' costs and in utility the individual elements of x_{jk} are perfect substitutes. This restriction implies that it is without loss of generality to assume that insurers choose g. Suppose that insurers did choose x to maximize:

$$\int_e \int_v MS(g(x), p, e, v) * [p - c(e, v, x)] dF(v|\sigma) dG(e))$$

with:

$$\begin{aligned} g(x) &= \beta' x \\ c(e, v, x) &= (\gamma_0 v + \mu_0 e + \psi_0) + (\gamma_1 v + \mu_1 e + \psi_1) * [\lambda' x] \end{aligned}$$

Claim 1. *If insurers are behaving optimally and for all m and n there exists plans k and k' such that $\frac{x_{kn}}{x_{km}} \neq \frac{x_{k'n}}{x_{k'm}}$ then $\frac{\beta_n}{\beta_m} = \frac{\lambda_n}{\lambda_m}$ for all n and m.*

Proof. Assume not. Then for some n and m, $\frac{\beta_n}{\beta_m} > \frac{\lambda_n}{\lambda_m}$. By assumption, there exist plans k and k' such that $\frac{x_{kn}}{x_{km}} < \frac{x_{k'n}}{x_{k'm}}$ (i.e. plans k and k' offer x_m and x_n in different ratios). If the characteristics in k' are optimally chosen, then the characteristics in k are not. Plan k can earn higher profits by increasing x_{kn} and decreasing x_{km}. Consider a one unit increase in x_{kn} and a $\frac{\beta_n}{\beta_m}$ decrease in x_{km}. This shift in characteristics leaves g_k unchanged. Therefore

$MS(g(x_k), p_k, e, v)$ is unchanged for all v. But profit margins will increase. The effect on costs is negative since:

$$\lambda_n \triangle x_{kn} + \lambda_m \triangle x_{km} =$$
$$\lambda_n - \lambda_m \frac{\beta_n}{\beta_m} < 0 \longleftrightarrow$$
$$\frac{\beta_n}{\beta_m} > \frac{\lambda_n}{\lambda_m}$$

The last line holds by assumption.

Thus, after normalization, insurers' cost functions are equivalent to $c(e, v, g) = (\gamma_0 v + \mu_0 e + \psi_0) + (\gamma_1 v + \mu_1 e + \psi_1) * g$. Insurers differentiate themselves via choices of overall insurance coverage, not the individual elements of x. It can be verified for each n and m that sufficient variation in product characteristics exists in the data so that the claim's assumption is satisfied.

C. Robustness to Multi-dimensional Types

Suppose that consumers' preferences for generosity are determined by their type $t_i \equiv (t_{i1}, t_{i2}, ..., t_{in})$ so that (after omitting irrelevant terms) utility is:

$$u_{ijk} = v_{gi}(t_{i1}, t_{i2}, ...t_{in}) * g_{jk} - p_{jk} + \xi_j^m + \varepsilon_{ijk}^m$$

Consumers' preferences for insurance now explicitly depend on their characteristics such as risk type, risk aversion, cognitive ability, etc. The restriction implicit in this specification is that insurers are unable to screen between consumers with the same v but different t. This assumption is admittedly strong. Once it is made, however, it is without loss of generality under certain assumptions to consider the more general, and realistic, cost function below:

$$\begin{aligned} c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply}) &= (\rho_{01} t_1 + ... + \rho_{0n} t_n + \psi_{jk}^0) + \\ &\quad (\rho_{11} t_1 + ... + \rho_{1n} t_n + \psi_{jk}^1) * g \end{aligned}$$

This allows individuals with the same preferences, but different characteristics, to impose different costs on insurers. Because an insurer can not discriminate between consumers with distinct t, but the same v, their cost function reduces to $E[c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply})|v]$. This can be seen by simply rearranging the insurer's profit function:

$$\int\limits_{t_1} ... \int\limits_{t_n} \sum_{k=1}^{n_j} MS_{jk}(g, p, v(t_{1,...,}t_n)|G, P, \Theta^{Demand})$$
$$*[S_m + p_{jk} - c_{jk}(t, g, \Theta^{Supply})]dH(t_1, ..t_n) =$$
$$\int\limits_{\widetilde{v}} \int\limits_{t:v(t)=\widetilde{v}} \sum_{k=1}^{n_j} MS_{jk}(g, p, \widetilde{v}|G, P, \Theta^{Demand})$$
$$*[S_m + p_{jk} - c_{jk}(t, g, \Theta^{Supply})]dH(t_1, ..t_n|v(t) = \widetilde{v})dF(\widetilde{v}) =$$
$$\int\limits_{\widetilde{v}} \sum_{k=1}^{n_j} MS_{jk}(g, p, \widetilde{v}|G, P, \Theta^{Demand}) * [S_m + p_{jk} - E[c_{jk}(t, g, \Theta^{Supply})|v(t) = \widetilde{v}]]dF(\widetilde{v})$$

Thus, ignoring the determinants of preferences for insurance can be made without loss of generality if:

$$\begin{aligned} E[c_{jk}(t, g, \Theta^{Supply})|v(t) &= \widetilde{v}] = c_{jk}(v, g, \Theta^{Supply}) = \\ &(\gamma_0 v + \psi_{jk}^0) + (\gamma_1 v + \psi_{jk}^1) * g \end{aligned}$$

where

$$\begin{aligned} E[c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply})|v] &= (\rho_{01}E[t_1|v] + ... + \rho_{0n}E[t_n|v] + \psi_{jk}^0) + \\ &(\rho_{11}E[t_1|v] + ... + \rho_{1n}E[t_n|v] + \psi_{jk}^1) * g \end{aligned}$$

Under specific conditions on the CDF of t, $H(t_1, ..t_n)$, this statement is true. Note that $c_{jk}(v.g, \Theta^{Supply})$ is linear in v and has an intercept $\psi_{jk}^0 + \psi_{jk}^1 * g$ when $v = 0$. Assume that $v(t) = 0$ if and only if $t_m = 0$ for each m. Thus, $E[c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply})|v]$ also has an intercept $\psi_{jk}^0 + \psi_{jk}^1 * g$ when $v = 0$. Now, note that $c_{jk}(v.g, \Theta^{Supply})$ is linear in v with slope $\gamma_0 + \gamma_1 * g$. $E[c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply})|v]$ is also linear in v if $\frac{dE[t_m|v]}{dv} = \triangle_m$ for all v and each m. For example, if $v(t_1, .., t_n) = t_1 + .. + t_n$ and $t_n\tilde{}\ |N(0, \sigma_{tn}^2)|$ for each n and $t_{n'}$ is uncorrelated with each t_n, this linearity condition holds. The slope of $E[c_{jk}(t_1, t_2, ..t_n, g, \Theta^{Supply})|v]$ in v is now equal to $(\rho_{01}\triangle_1 + ... + \rho_{0n}\triangle_n) + (\rho_{11}\triangle_1 + ... + \rho_{1n}\triangle_n) * g$. Thus, the restricted model, with $c_{jk}(v.g, \Theta^{Supply})$ can duplicate the more general model if $\gamma_0 = \rho_{01}\triangle_1 + ... + \rho_{0n}\triangle_n$ and $\gamma_1 = \rho_{11}\triangle_1 + ... + \rho_{1n}\triangle_n$.

D. Estimation

Demand.—

Construction of Demand Moments Under the assumption that the horizontal shock to preferences, ε_{ijk}^{m}, is a type-1 logit error, the probability that individual i enrolls in $j'k'$ in market m, conditional on $i's$ unobserved type, v_i, can be written as:

$$\begin{aligned} Prob(j'k'|e_i, v_i, m, \Theta) &= \frac{e^{\theta_i g_{j'k'} - \alpha_m p_{j'k'} + \xi_{j'm}}}{e^{\lambda_0 e_i + \chi_m} + \sum\limits_{j \in \Omega_m} \sum\limits_{k \in \Omega_{mj}} e^{\theta_i g_{jk} - \alpha_m p_{jk} + \xi_{jm}}} \\ &= \frac{e^{u_{ij'k'} + \xi_{j'}}}{e^{u_{i0}} + \sum\limits_{j \in \Omega_m} \sum\limits_{k \in \Omega_{mj}} e^{u_{ijk} + \xi_j}} \end{aligned}$$

where Ω_m is the set of firms that have entered market m and Ω_{jm} is the set of products offered by each firm j that is participating in market m.

Three other probability functions are used in estimation. The probability that consumer i enrolls in a product offered by firm j' in market m can be written by summing over $Prob(j'k'|v_i)$ for each $k' \in \Omega_{jm}$:

$$Prob(j'|e_i, v_i, m, \Theta) = \frac{\sum\limits_{k' \in \Omega_{j'm}} e^{u_{ij'k'} + \xi_{j'}}}{e^{u_{i0}} + \sum\limits_{j \in \Omega_m} \sum\limits_{k \in \Omega_{mj}} e^{u_{ijk} + \xi_j}}$$

The probability of i enrolling in product k' conditional on enrolling in firm j' can be written as:

$$Prob(k'|j', e_i, v_i, m, \Theta) = \frac{e^{u_{ij'k'}}}{\sum\limits_{k \in \Omega_{mj'}} e^{u_{ijk}}}$$

Note that firm level $\xi_{j'}$ falls out of $Prob(k'|j', v_i, m, \Theta)$.

The probability of i enrolling in product $j'k'$ in market m', and then product $\widetilde{\widetilde{jk}}$ one year later in market $\widetilde{m}$ can be written as:

$$Prob(j'k', \widetilde{\widetilde{jk}}|e_i, v_i, m', \widetilde{m}, \Theta) = \frac{e^{u_{ij'k'} + \xi_{j'}}}{1 + \sum\limits_{j \in \Omega_m} \sum\limits_{k \in \Omega_{mj}} e^{u_{ijk} + \xi_j}} * \frac{e^{u_{i\widetilde{\widetilde{jk}}} + \xi_{\widetilde{j}'}}}{1 + \sum\limits_{j \in \Omega_{\widetilde{m}}} \sum\limits_{k \in \Omega_{\widetilde{m}j}} e^{u_{ijk} + \xi_j}}$$

It is necessary to work with the unconditional share probabilities. These are obtained simply by integrating over $Prob(j'k'|e_i, v_i, m, \Theta)$ and $Prob(j|e_i, v_i, m, \Theta)$ with respect to e and v:

$$Prob(j'k'|m,\Theta) = \int_v \int_e \frac{e^{u_{ij'k'}+\xi_{j'}}}{e^{u_{i0}} + \sum_{j\in\Omega_m} \sum_{k\in\Omega_{mj}} e^{u_{ijk}+\xi_j}} dF(v|\sigma)dG(e)$$

$$Prob(j'|m,\Theta) = \int_v \int_e \frac{\sum_{k'\in\Omega_{j'm}} e^{u_{ij'k'}+\xi_{j'}}}{e^{u_{0i}} + \sum_{j\in\Omega_m} \sum_{k\in\Omega_{mj}} e^{u_{ijk}+\xi_j}} dF(v|\sigma)dG(e)$$

$$\begin{aligned} Prob(k'|j',m,\Theta) &= \int_v \int_e \frac{e^{u_{ij'k'}+\xi_{j'}}}{\sum_{k\in\Omega_{mj'}} e^{u_{ijk}+\xi_j}} dF(v|j',\sigma)dG(e) \\ &= \frac{Prob(j'k'|m,\Theta)}{Prob(j'|m,\Theta)} \end{aligned}$$

$$Prob(j'k',\widetilde{\widetilde{jk}}|m,\Theta) = \int_v \int_e \frac{e^{u_{ij'k'}+\xi_{j'}}}{e^{u_{i0}} + \sum_{j\in\Omega_m} \sum_{k\in\Omega_{mj}} e^{u_{ijk}+\xi_j}} * \frac{e^{u_{i\widetilde{j}\widetilde{k}}+\xi_{\widetilde{j}'}}}{e^{u_{i0}} + \sum_{j\in\Omega_{\widetilde{m}}} \sum_{k\in\Omega_{\widetilde{m}j}} e^{u_{ijk}+\xi_j}} dF(v|\sigma)dG(e)$$

With a set of L random draws from $f(v|\sigma)$ and $g(e)$ I simulate the above probabilities with $\widehat{Prob}(j'k'|m,\Theta) = \frac{1}{L}\sum_{l=1}^{L} \frac{e^{u_{ij'k'}+\xi_{j'}}}{e^{u_{i0}}+\sum_{j\in\Omega_m}\sum_{k\in\Omega_{mj}} e^{u_{ijk}+\xi_j}}$, $\widehat{Prob}(j'|m,\Theta) = \frac{1}{L}\sum_{l=1}^{L} \frac{\sum_{k'\in\Omega_{j'm}} e^{u_{ij'k'}+\xi_{j'}}}{e^{v_{0i}}+\sum_{j\in\Omega_m}\sum_{k\in\Omega_{mj}} e^{u_{ijk}+\xi_j}}$, and $\widehat{Prob}(k'|j',m,\Theta) = \frac{\widehat{Prob}(j'k'|m,\Theta)}{\widehat{Prob}(j'|m,\Theta)}$. All moments are constructed using these simulated probabilities.

Three forms of moment conditions are used in estimation. The population moments, described in the main text are repeated below:

(16)
$$E[x_{jk} * (d_{ijk} - Prob(k|j,m,\Theta))|ij] = 0$$

(17)
$$E[x_{jkl} * x'_{j'k'} * (d_{ijk,j'k'} - Prob(jk\&j'k'|m,m',\Theta))|i] = 0$$

$$E[f(x_{-jk}|x_{jk}) * (d_{ijk} - Prob(k|j,m,\Theta))|ij] = 0 \tag{18}$$

The corresponding sample moments are:

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{x_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta) - d_{ijk}] = 0$$

$$\frac{1}{M}\sum_{i=1}^{M}\frac{1}{|\Omega_m| * |\Omega_{m'}|}\sum_{jk\in\Omega_m}\sum_{j'k'\in\Omega_{m'}} x_{jk}x'_{j'k'} * (Prob(jk, j'k'|m,\Theta) - d_{i,jk,j'k'}) = 0$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{f(x_{-jk}|x_{jk})}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta) - d_{ijk}] = 0$$

Above, j sums over insurers, k over products and i over MCBS respondents. These summations are divided through by $\sqrt{n_{jm}}$ instead of n_{jm} to ensure that the variance of residuals interacted with instruments ($\frac{1}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta) - d_{ijk}]$) does not depend on the number of survey respondents from each firm. This leads to more weight being placed on contracts for whom we observe more individuals enrolling.

Accounting for Simulation Error Consider a sample moment condition employed in demand:

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta) - d_{ijk}] = 0$$

where $Prob(k|j,m,\Theta)$ is the true probability of enrolling in product k conditional on enrolling in firm j. Rewrite this moment condition, using $\widehat{Prob}_m(k|j,m,\Theta)$ to represent simulated probabilities:

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)$$
$$-\widehat{Prob}(k|j,m,\Theta)+\widehat{Prob}(k|j,m,\Theta)-d_{ijk}]=$$
$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-\widehat{Prob}(k|j,m,\Theta)]$$
$$+\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[\widehat{Prob}(k|j,m,\Theta)-d_{ijk}]=0$$

Because simulation error is independent of the randomness inherent to $\sum_i[\widehat{Prob}(k|j,m,\Theta)-d_{ijk}]$, the variance of the sample moment is equal to:

$$Var\left[\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-\widehat{Prob}(k|j,m,\Theta)]\right]+$$
$$Var\left[\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[\widehat{Prob}(k|j,m,\Theta)-d_{ijk}]\right]$$

The first term is variance due to simulation error. Note that $\widehat{Prob}_m(k|j,m,\Theta)]$ is the sum of $nsim$ draws of $Prob(k|j,v_i,m,\Theta)]$ where v_i is distributed according to $F(v|\sigma)$. Because $E\left(Prob(k|j,v_i,m,\Theta)]\right)=Prob(k|j,m,\Theta)$, the delta method implies:

$$\widehat{Prob}_m(k|j,v_i,m,\Theta)]-Prob(k|j,m,\Theta)\tilde{}N(0,\frac{\partial\widehat{Prob}_m(k|j,v_i,m,\Theta)]}{\partial v}\sigma\frac{\partial\widehat{Prob}_m(k|j,v_i,m,\Theta)]}{\partial v'})$$

Using this result, it is straightforward to construct the variance of $\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-\widehat{Prob}_m(k|j,m,\Theta)]$. Analogous procedures can be used for the remaining moments.

Selection in Demand Moments Consider a moment condition that interacts Z_{jk} with the random residuals, $\sum_i[Prob(k|j,m,\Theta)-d_{ijk}]$. Assume for a particular firm j that $Z_{jk}=Z_j$ for all k. That firm's contribution to the sample moment condition is equal to zero, regardless of the choice of parameters:

$$\begin{aligned}
\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{Z_j}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-d_{ijk}] &= \\
\frac{Z_j\sqrt{n_{jm}}}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{1}{n_{jm}}\sum_i[Prob(k|j,m,\Theta)-d_{ijk}] &= \\
\frac{Z_j}{|\Omega_{jm}|}\frac{1}{\sqrt{n_{jm}}}\sum_{k=1}^{|\Omega_{jm}|}(Prob(k|j,m,\Theta)-\frac{1}{n_{jm}}\sum_i d_{ijk}) &= \\
\frac{Z_j}{|\Omega_{jm}|}\frac{1}{\sqrt{n_{jm}}}*[\sum_{k=1}^{|\Omega_{jm}|}(Prob(k|j,m,\Theta)-\sum_{k=1}^{|\Omega_{jm}|}\frac{1}{n_{jm}}\sum_i d_{ijk}] &= \\
\frac{Z_j}{|\Omega_{jm}|}\frac{1}{\sqrt{n_{jm}}}*[1-1] &= 0
\end{aligned}$$

To model this selection issue, let s_{jZ} be an indicator variable that equals one if and only if there is variation in Z_{jk} across k. Interacting s_{jZ} with Z and yields the following moment condition:

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\frac{1}{|\Omega_{jm}|}\sum_{k=1}^{|\Omega_{jm}|}\frac{s_{jZ}*Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-d_{ijk}]$$

This sample moment condition should equal zero in expectation only if for each firm,

$$E\left[\sum_{k=1}^{|\Omega_{jm}|}\frac{s_{jZ}*Z_{jk}}{\sqrt{n_{jm}}}\sum_i[Prob(k|j,m,\Theta)-d_{ijk}]\right]=0$$

The discussion in Wooldridge (pg 573), argues that the selection case here is of the most favorable kind. Note that s_{jZ} is a deterministic function of the instruments themselves, and therefore $E\left[s_{jZ}(Z)*Z_{jk}*\epsilon_{jk}\right]=0$ if $E\left[\epsilon_{jk}\mid Z\right]=0$ via iterated expectations. That is indeed the case, and therefore, the sample moment conditions still equal zero in expectation.

Supply.—Recall that in each market the collection of plan characteristics, $\{g_{jk},p_{jk}\}_{jk\in\Omega_m}$ satisfy a Nash Equilibrium. Therefore, all firms' first order conditions with respect premiums and insurance generosity must equal zero:

$$\frac{\partial Profit_j}{\partial g_{jk}}=\frac{\partial Profit_j}{\partial p_{jk}}=0\ \forall\ j,k$$

Backing out cost shocks ψ Firms' first order conditions for p_{jk} and g_{jk} reduce to the following set of linear equations. First, for the first order condition on premium:

$$
\begin{aligned}
\frac{\partial Profit_{jkm}}{\partial p_{jk'}} &= c^0_{p_{jk'}} + \sum_{k=1}^{m_j}\sum_{z=0}^{2} c^{k,z}_{p_{jk'}} * \psi^z_{jk} = 0 \\
&\quad where \\
c^0_{p_{jk'}} &= \int\limits_e \int\limits_\nu \sum_{k=1}^{m_j} \frac{\partial MS_{jk}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial p_{jk'}} * \\
&\quad [S_m + p_{jk} - \gamma_0' * \nu - \mu_0 * e] \\
&\quad -(\gamma_1' * \nu + \mu_1 * e) * g_{jk}] dF(\nu|\sigma)dG(e) \\
&\quad + \int\limits_e \int\limits_\nu MS_{jk'}(g,p,e,\nu|G,P,\Theta^{Demand}) dF(\nu|\sigma)dG(e) \\
c^{k,0}_{p_{jk'}} &= -\int\limits_e \int\limits_\nu \frac{\partial MS_{jk}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial p_{jk'}} dF(\nu|\sigma)dG(e) \\
c^{k,1}_{p_{jk'}} &= -g_{jk} \int\limits_e \int\limits_\nu \frac{\partial MS_{jk}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial p_{jk'}} dF(\nu|\sigma)dG(e)
\end{aligned}
$$

and below is the first order condition on quality:

$$
\begin{aligned}
\frac{\partial Profit_{jkm}}{\partial g_{jk'}} &= c^0_{g_{jk'}} + \sum_{k=1}^{m_j}\sum_{z=0}^{1} c^{k,z}_{g_{jk'1}} * \psi^1_{jk'} = 0 \\
&\quad where \\
c^0_{g_{jk'1}} &= \int\limits_e \int\limits_\nu \sum_{k=1}^{m_j} \frac{\partial MS_{jk}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial g_{jk'1}} * \\
&\quad [S_m + p_{jk} - \gamma_0'\nu - \mu_0 e - (\gamma_1'\nu + \mu_1 e) g_{jk'}] dF(\nu|\sigma)dG(e) \\
&\quad - \int\limits_e \int\limits_\nu MS_{jk'}(g,p,e,\nu|G,P,\Theta^{Demand}) * (\pi_1 + \gamma_1' * \nu + \delta_1' * \ Z_{jm}) dF(\nu|\sigma)dG(e) \\
c^{k,1}_{g_{jk'1}} &= -g_{jk} \int\limits_e \int\limits_\nu \frac{\partial MS_{jk}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial g_{jk'}} dF(\nu|\sigma)dG(e) \ \ \forall\ k \neq k' \\
c^{k',1}_{g_{jk'1}} &= -g_{jk'} \int\limits_e \int\limits_\nu \frac{\partial MS_{jk'}(g,p,e,\nu|G,P,\Theta^{Demand})}{\partial g_{jk'}} dF(\nu|\sigma)dG(e) - \\
&\quad \int\limits_\nu MS_{jk'}(g,p,e,\nu|G,P,\Theta^{Demand}) dF(\nu|\sigma)dG(e)
\end{aligned}
$$

By inspection, these linear equations are non-redundant. Backing out unobserved cost shocks is done via the following matrix algebra:

$$\begin{pmatrix} \widetilde{\psi}^0_{j1} \\ \widetilde{\psi}^1_{j1} \\ . \\ . \\ \widetilde{\psi}^0_{jm_j} \\ \widetilde{\psi}^1_{jm_j} \end{pmatrix} = -1 * \begin{pmatrix} c^{1,0}_{p_{j1}} & c^{1,1}_{p_{j1}} & . & . & c^{m_j,0}_{p_{j1}} & c^{m_j,0}_{p_{j1}} \\ c^{1,0}_{g_{j1}} & c^{1,1}_{g_{j1}} & . & . & c^{m_j,0}_{g_{j1}} & c^{m_j,1}_{g_{j1}} \\ . & . & . & . & . & . \\ . & . & . & . & . & . \\ c^{1,0}_{p_{jm_j}} & c^{1,1}_{p_{jm_j}} & . & . & c^{m_j,0}_{p_{jm_j}} & c^{m_j,1}_{p_{jm_j}} \\ c^{1,0}_{g_{jm_j}} & c^{1,1}_{g_{jm_j}} & . & . & c^{m_j,0}_{g_{jm_j}} & c^{m_j,1}_{g_{jm_j}} \end{pmatrix}^{-1} * \begin{pmatrix} c^0_{p_{j1}} \\ c^0_{g_{j1}} \\ . \\ . \\ c^0_{p_{jm_j}} \\ c^0_{g_{jm_j}} \end{pmatrix}$$

Construction of Supply Side Moments Supply side moments are constructed by exploiting assumed population moment conditions relating firms' unobserved costs shocks to instruments. These conditions are:

$$\begin{aligned} E[\psi^n_{jk}] &= 0 \ \forall \ jkn \\ E[X_m * \psi^n_{jk}] &= 0 \ \forall \ jkn \\ E[Z_j * \psi^n_{jk}] &= 0 \ \forall \ jkn \end{aligned}$$

$$\begin{aligned} E[f(Z'_{-jm}) * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[N_{Prod,m} * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[N_{Firm,m} * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[f(Z'_{-jm}) * N_{Prod,m} * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[f(Z'_{-jm}) * N_{Firm,m} * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[\xi_j * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \\ E[f(\xi_{-j}) * \psi^n_{jk}] &= 0 \quad \forall \ j,k,n \end{aligned}$$

The corresponding sample moment conditions are below. Notation is the same as above. Within each firm, the residual is constructed as $\frac{1}{\sqrt{m_j}} \sum_{k=1}^{m_j} \psi^n_{jk}$ so that the variance does not depend on the number of products offered by firm j, m_j.

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(\frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(X_m' * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(Z_{jm}' * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(f(Z_{-jm}') * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(N_{\Pr od,m} * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(N_{Firm,m} * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(f(Z_{-jm}') * N_{\Pr od,m} * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(f(Z_{-jm}') * N_{Firm,m} * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(\xi_j * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

$$\frac{1}{N_j}\sum_{j=1}^{N_j}\left(f(\xi_{-j}) * \frac{1}{\sqrt{m_j}}\sum_{k=1}^{m_j}\psi_{jk}^n\right) = 0 \quad \forall\ j,k,n=0,1$$

Simulation Error in Supply Moments Simulation error is also present in the supply moments because they are constructed from cost shocks ψ which are non-linear functions of the simulated market share functions. Because the simulation error enters each moment condition in a non-fashion, biases may be present. Changing the number of simulation draws, however, does not have large effects on parameter estimates. To correct variance estimates for this simulation error, I repeatedly take samples of simulation draws to simulate the variation of each moment condition due to simulation. As above, this variation is independent of variation in the sample moments due to observable data, and in practice quite small. All reported standard errors are corrected for this simulation error.

E. Summary of Estimation Routine

The estimation routine is simulated method of moments. It searches over $\Theta \equiv (\alpha, \beta, \lambda, \theta, \sigma, \gamma, \mu)$ for the minimum of a criterion function, $\Psi(\Theta) = M(\Theta)' * W * M(\Theta)$, where $M(\Theta)$ consists of sample analogs of the demand and supply moments described above. These sample analogs (a full list is provided in the appendix) are formed summing over markets, plans, and MCBS respondents.

Evaluating Ψ can be broken into four steps for each choice of Θ :

1. Use the BLP contraction mapping to generate $\widehat{\xi}$
2. Plug $\widehat{\xi}$ into the insurers' first order conditions and back out $\widetilde{\psi}$.
3. Regress $\widetilde{\psi}$ on a constant, Z_{jm} to find $\widehat{\pi}$, $\widehat{\delta}$, and ψ
4. Evaluate the remaining supply and demand moment conditions and calculate the value of Ψ at Θ.

The market share probabilities and insurers' cost shocks ψ are continuous in the parameters. In the text above, intuition for identification was given. Assuming standard regularity conditions yields consistency of the parameter estimates and their asymptotic distribution (Newey & McFadden 1994).

In practice, parameter estimates were found using genetic search and Nelder-Mead alogirthms provided by Matlab. Initial parameter estimates were found by estimating demand and supply separately and then, simultaenously.

F. Properties of $\widehat{\Theta}$

Following Newey & McFadden (1994) the following conditions are necessary for $\widehat{\Theta} \xrightarrow{p} \Theta_0$, where $\widehat{\Theta}$ minimizes $\Omega_n(m, \Theta) \equiv (M_n(\Theta)' * W * M_n(m, \Theta)$.

1. $\Omega(m, \Theta)$ is uniquely minimized at Θ_0
2. The parameter space Θ is compact.
3. $\Omega(m, \Theta)$ is continuous
4. $\Omega_n(m, \Theta)$ converges uniformly in probability to $\Omega(m, \Theta)$.

The first two conditions are assumed. $M_n(\Theta)$ is highly non-linear, thus proving identification is a challenge. Intuition for identification, however, is provided throughout the main text. Condition 3 is satisfied by simple inspection of moment conditions. All market share functions entering demand moments are continuous functions over the entire parameter space Θ. Similarly, firms' first order conditions are all continuously differentiable at all Θ. Thus, the unobserved cost shocks, ψ are continuous functions of Θ since inversion preserves continuity. Because $M_n(\Theta)$ is continuous for all Θ, so is $\mathrm{M}_0(m, \Theta)$ and $\Omega(m, \Theta)$. The fourth condition is true under standard regularity conditions.

G. Additional Tables

Table A1: Estimated regression coefficients for regression of Medical Expenditures on Observable Characteristics

High School Degree	0.1159	Ever Smoked Cigarettes	0.053
	(0.0192)**		(0.0159)**
Some College	0.19186	Flu Shot Last Year	0.1645
	(0.02121)**		(0.0166)**
BA Degree	0.2259	Prostate Exam Last Year	0.3444
	(0.0252)**		(0.0225)**
Widowed	0.0626	Mammogram Last Year	0.2651
	(0.0186)**		(0.0216)**
Divorced	0.0714	Pap Smear Last Year	0.1595
	(0.0297)*		(0.0225)**
Never Married	-0.0839	History with Cancer	0.4227
	(0.0425)*		(0.0184)**
Income 10-20K	-0.0351	History with Diabetes	0.3901
	(0.0210)		(0.0188)**
Income 20-30K	0.0289	History of High Blood Pressure	0.1255
	(0.0241)		(0.0154)**
Income 30-40K	0.0698	Previous Heart Attack	0.3548
	(0.0290)*		(0.0203)**
Income 40-50K	0.1473	History of Arthritis	0.1068
	(0.0339)**		(0.0238)**
Income > 50K	0.1916	History of Rheumatoid Arthritis	0.1273
	(0.0318)**		(0.0155)**
Difficulty Walking	0.2534	History of Psychological Problems	0.1441
	(0.0200)**		(0.0231)**
Difficulty Stooping	0.0645	Body Mass Index	-0.0045
	(0.0178)**		(0.0011)**
Difficulty Lifting Items > 10 lbs	0.1632	History of Emphysema	0.2558
	(0.0209)**		(0.0210)**
Difficulty Social Activities	0.519	History of Osteoporosis	0.1195
	(0.0245)**		(0.0199)**
Difficulty with Memory	-0.0063	Health Perceived Excellent	-0.9854
	(0.0245)		(0.0397)**
Difficulty handling Money	-0.0049	Health Perceived Very Good	-0.7642
	(0.0364)		(0.0366)**
Difficulty using Phone	-0.1534	Health Perceived Good	-0.4735
	(0.0296)**		(0.0346)**
Difficulty Bathing	0.2996	Health Perceived Fair	-0.3036
	(0.0259)**		(0.0339)**
Difficulty with Housework	0.0617		
	(0.0196)**		
Difficulty preparing Meals	0.0848		
	(0.0337)*		
Observations		49177	
R-squared		0.2017	

Notes: Dependent variable is the natural log of total medical expenditures covered by parts A and B of Medicare. Regression is limited to enrollees in traditional Medicare. Regression includes year, race, and sex*age fixed effects. Standard errors in parentheses. * indicates significant at 5% level. ** indicates significant at 1%

Table A2: Reduced Form Supply Side Regressions

	Generosity	Premium
Constant	2.413	90.91
	(.163)***	(7.68)***
Market Controls		
AB Rate / 10	0.02	0.469
	(.004)***	(.17)**
AB Rate 1997 / 10	0.001	-1.34
	(.003)	(.134)***
Population / 1000	(.003)	0.012
	(.0005)***	(.021)
Per Capita Income / 1000	0.102	11.38
	(.035)**	(1.48)**
# Hospitals	-0.014	-0.429
	(.004)***	(.163)**
2001 Dummy	-0.54	2.76
	(.048)***	(2.04)
2002 Dummy	-1.178	-1.13
	(.059)***	(2.49)
2003 Dummy	-1.066	14.86
	(.051)***	(2.18)***
Insurer Controls		
HMO Dummy	-0.236	-1.99
	(.051)***	(3.26)
Total Assets ($ million)	0.04	-1.5
	(.027)	(1.16)
Net Income ($ million)	0.001	-0.026
	(.001)	(.023)
# Network Physicians / Enrollee	-0.021	0.788
	(.004)***	(.188)***
% Administrative Expenses	0.053	-11.86
	(.022)*	(1.33)***
Medical Expenditures / Enrollee	0.0006	-0.004
	(.0002)***	(.007)
% Business Medicare	0.014	-2.9909
	(.011)	(.4601)***
N	5622	5622
R-Squared	0.13	0.292

Notes: Firm controls lagged one year. Standard Errors are in parentheses. *,**,and *** indicate significance at the 5%, 1% and 0.1% levels