© *The Journal of Risk and Insurance*, 2007, Vol. 74, No. 1, 115-140

The Demand for Health Insurance in the Group Setting: Can You Always Get What You Want?

Mark V. Pauly
Bradley Herring

Abstract

To what extent do health benefits obtained in the employment-based setting reflect individual preferences? We examine this question by comparing the relationship between person-level characteristics and the plans they obtain in a group setting to the relationship observed in the individual insurance market, using data from the 1996-1997 and 1998-1999 Community Tracking Study's Household Surveys. We also examine the effect of unions on group choice. Our structural models of the demand for insurance indicate that plans obtained in the group setting often reflect underlying individual preferences for insurance, but we consistently observe significantly different effects of ethnicity and unionization.

Introduction

The great bulk of private health insurance in the United States is obtained in the employment setting. Employers (sometimes in concert with unions) choose whether to offer insurance, what plan to offer, and how many plans to offer. However, in highly competitive labor markets that typically characterize the American economy, employees usually have a choice among jobs. At least to some extent, that job choice is influenced by the attractiveness of the combination of money wages and health benefits that a firm offers. Indeed, the primary reason why a profit-maximizing firm would offer costly health benefits is because workers of a given quality level would then be willing to work for money wages reduced to a sufficient extent to cover the cost of the benefit. The fundamental modeling (and, we will argue, policy) question is how closely employers are able to match their offerings with employee demand.

Mark V. Pauly is at the Health Care Systems Department, The Wharton School, University of Pennsylvania, 3641 Locust Walk, Philadelphia, PA 19104 and Bradley Herring is at the Department of Health Policy & Management, Emory University's Rollins School of Public Health, 1518 Clifton Road NE, Atlanta, GA 30322. The authors can be contacted via e-mail: pauly@wharton.upenn.edu and bradley.herring@emory.edu. The authors thank the Robert Wood Johnson Foundation's Changes in Health Care Financing and Organization (HCFO) Initiative for financial support for this work. The authors appreciate helpful comments from seminar participants at the 2003 NBER Health Care Program Meeting, Jose Pagan, and two anonymous referees.

At one extreme, there is the Tiebout (1956) model of perfect matching, originally used to explain how households choose communities with different combinations of public services and taxes. Goldstein and Pauly (1976) demonstrate that, in such a model, if there are many firms offering different levels and types of health benefits that employees may choose among, and if worker mobility is nearly costless, the equilibrium will be one in which each worker will be matched with a firm offering the health insurance policy the worker would most prefer, given its cost (net of any tax subsidies). They then raise the question, still not yet resolved, of how closely actual group insurance provision matches this benchmark case. In this model, there are two important policy implications: workers get the health insurance they want (i.e., "perfect matching"), and workers pay the (gross) cost of that insurance in the form of lower wages (i.e., "perfect incidence").

Given the difficulty of directly estimating the incidence of health benefits on wages empirically, there is therefore an important byproduct from this model: an observation that workers get the health insurance they want, given its true cost, would be consistent with the hypothesis of perfect incidence of health benefits on wages. If workers who would demand more insurance at a given price actually get it in the group setting, they are behaving as if they actually paid the price. For example, if higher income workers get more coverage than lower income workers, it is plausible that the higher income workers are willing to sacrifice more wages to get it. The argument is not wholly conclusive, because we do not know how decisions on coverage might be made if workers are not paying for all of it, but were rather paying a constant proportion of the cost. However, if the incidence was on employer profits (rather than wages), there is no obvious reason why employers would provide more coverage when workers have higher incomes. Therefore, testing demand matching is an indirect way of determining incidence.

Note that, in this Tiebout model, many individual firms might choose their insurance offerings in many different ways. For example, firms might just pick benefit offerings randomly, but the workers might then adjust by moving across firms. Or firms might choose offerings to match the preferences of the marginal workers they want to attract and retain. "Unionized" firms might choose the coverage desired by the median union member or perhaps the older long-term union members. Regardless of the process by which equilibrium is reached, a fundamental question is whether unsatisfied or excess demand exists at the margin or, alternatively, whether some employers will not offer plans when workers demand them. "Perfect matching" is a property of the system in equilibrium, and we need not be concerned with the behavior of individual firms (in terms of specifically *how* that equilibrium is reached).

At the other extreme, workers might be assumed to be assigned to firms and jobs for life. In this case the only reason to offer health insurance is because health insurance may affect health, which in turn affects worker productivity. In this model, employers might choose plans that employees dislike; this would be most consistent with the backlash against HMOs or other managed care plans chosen by employers. Even the potential tax benefits to employees would at least be a matter of indifference to an employer seeking to minimize its total compensation costs, since employees are assumed not to move in response to benefits.

Of course, we know that neither of these models is exactly accurate. Labor mobility is not perfect, but it is not nonexistent either, so probably matching is not perfect, but also exists to some extent. For instance, Monheit and Vistnes (1999) find that workers with weak preferences for health insurance sort themselves into jobs (such as those at small firms) that do not offer coverage. Both Moran, Chernew, and Hirth (2001) and Bundorf (2002) find that the diversity of workers within a firm increases the breadth of plans offered to workers, but fixed administrative costs limit the extent to which multiple plans may be offered to match all worker preferences. However, Blumberg and Nichols (2004) and Blumberg and Cancian (2004) offer theory and some evidence that not all workers have available to them roughly equivalent job options with and without insurance. They find that in many small labor markets some job-health benefit combinations do not exist. These analyses do not address the question of how close the potential matching is. Put slightly differently, these papers do not offer reasons why any large number of workers would be unable to get a plan close to what they want. The important empirical question then is: how close are actual group health insurance markets to either model presented above? That is the question we consider in this paper.

THE BENCHMARK MODEL AND THE TEST

The strategy we follow in this analysis is simple. We first use data from the individual, or "nongroup", health insurance market to estimate demand for insurance in that market, and we assume that this is the "true" or unbiased demand curve for all persons. Of course, the price of insurance is generally (though not always) higher in the individual market than in the group market, but that fact does not limit our ability to estimate a demand curve, with insurance demand defined as a function of the demographic characteristics of buyers.

We then use this individual demand curve as a benchmark to compare with the results from the group insurance process. For persons defined by a given set of characteristics, we ask whether the group insurance they get is the same as what they would have been predicted to demand in the individual market. If, for example, (ignoring interactions) we find that high-income individuals are 50 percent less likely than low-income individuals to demand a managed care plan (rather than an indemnity plan) in the individual market, we look to see whether the odds of getting a managed care plan for people at similar income levels differ by the same amount in the group market. In effect, we compare the plan differences "predicted" from the individual market demand curve for people actually in the group market with the insurance that they actually get. A finding that the variation in insurance the person ends up with is approximately the same would be evidence consistent with the view that the employment-based system is close to the Tiebout model of perfect matching. Implications about incidence and the absence of need to correct or override what employers and employees are doing would then follow. Obviously we are assuming that there are no strong interactions between these characteristics and the loading for insurance.

The data available to us (and any other data of which we are aware) does not list all conceivable health plan characteristics. Hence, we will be limited to comparing those characteristics for which we have reliable measures. Formally, we first obtain the "benchmark" insurance demand curve for a sample of people in the individual insurance market, by estimating several selected measures of the quantity or quality of insurance demanded as a function of a basic set of person-level demographic characteristics. We then estimate the demand curve for a sample of people in the group market using the same basic set of person-level demographic characteristics.

The most important statistical test is whether there is a significant difference between the vector of coefficients estimated in the individual market model and those estimated in the group market analog. If the hypothesis of similarity is rejected, we then examine coefficients on specific variables to determine which ones differ significantly and to what extent. (For this reason, we choose a simple specification for each characteristic so that our test for the difference across samples for a particular characteristic is stronger.) Another relevant test for the degree of "matching" concerns the residuals from either regression. It is possible that all the coefficients might be the same, but some measure of goodness of fit, such as the R-squared, might differ. For instance, if we found that the R-squared was considerably higher for the model of demand in the individual market than for the model of demand in the group market, even if the estimated coefficients were the same, we still might conclude that some unmeasured influences have different effects in the group setting relative to the individual setting.

Other Considerations

Our approach of comparing the insurance obtained in the individual and group markets is somewhat similar to that of Browne (1992), but we extend his work in two primary ways. One extension is that we consider several aspects of insurance demand besides the amount of cost-sharing among those who obtain insurance; we consider, in addition, both the extent of managed care restrictions placed on care and whether any insurance is obtained. Another extension of Browne's work is that we examine whether differences between individual and group demand exist across a whole set of demographic characteristics; his paper focused on the effect of expected medical expense on adverse selection.

Like Browne's work, though, we face a similar limitation in interpreting results for effects of health status on the demand for insurance. He found that people with lower expected expense had plans with more cost-sharing in the individual market (relative to the plans low-risk people obtained in the group market), and concluded that this finding is consistent with more adverse selection existing in the individual market. However, he could not rule out that his finding was "explained by low risks being off their demand curve in the group market" (p. 13). As such, our "benchmark" insurance demand curve using people in the individual market incorporates not only the underlying demand for insurance but also any differential impact of adverse selection or medical underwriting on the insurance obtained.

Although the theoretical implication of the Rothschild-Stiglitz (1976) model is that when an equilibrium exists, low risks will obtain a lower quantity of insurance, it is unclear how much (if any) adverse selection actually exists in the health insurance market (Cardon and Hendel, 2001). Probably selection (adverse or taste-related) would be greater in the individual market, although the spread (since Browne's work) of multiple-choice group insurance with no risk rating of employee-paid premiums makes even this prediction somewhat less uncertain. Medical underwriting by individual and small-group insurers to exclude issuing coverage to some (but certainly not all) high-risk people is also thought to exist (particularly when state rating regulations are in effect), but the extent to which underwriting actually occurs in the individual and small group markets is not clear either. We must therefore rely on our empirical results, but even here interpretation will be hard.

However, the fact that these "other" things affect the demand for insurance in the individual market is not necessarily problematic for our analysis, so long as we keep in mind that the fundamental question we are asking in this analysis is whether the process of obtaining health insurance through groups somehow "shifts" the demand curve away from the situation that would prevail if people obtained their insurance individually (given the way that the individual market actually functions) and interpret our results accordingly. Similarly, we do not incorporate federal or state insurance regulations explicitly in our analysis, but we do comment in several areas below on how such regulations may affect the results we observe for relevant demographic characteristics.

The Effect of Unionization

A more complex question concerns the interpretation of any impact of *additional* group characteristics on the insurance people receive. For instance, one might observe that individuals of similar characteristics are more likely to be insured or have a particular type of insurance if they work for a unionized firm or a firm with a large proportion of high-wage workers. Clearly, if no measured group characteristic is significantly associated with group demand, or if firm characteristics have the same effect on the demands of workers who buy individually as those who buy in groups, we would conclude that the process of obtaining insurance through a group makes no difference. But, for example, what if people who work for unionized firms are more likely to have a particular kind or level of coverage?

Such a finding would surely be consistent with a hypothesis that unionization affects group insurance demand, a hypothesis suggested by Goldstein and Pauly (1976). They explored a number of possible models of choice of plan in unionized firms. Although one model was a simple median voter story, they also hypothesized that union choice may be dominated by the interests of older, long-term members who would generally be higher risk and interested in more generous coverage. A positive impact of unionization on plan generosity is consistent with this latter model.

But might evidence for a unionization effect at the firm level also be consistent with a view that unionization per se makes no difference, but rather that people with strong demands for insurance also choose to work for firms with unions? The problem is that the firm or group characteristics attached to a given person may be endogenous. How can we deal with this? It is a self-selection problem, but there are no obvious identifying variables at the firm or worker level. Rather than characterize a worker by the unionization status of his firm (which is likely to be endogenous), we use a measure of the extent of unionization in the worker's industry in the region. This measure is more plausibly regarded as exogenous, since it is unlikely that workers choose their job skills or where to live based upon their demand for health insurance.

ESTIMATING THE INDIVIDUAL DEMAND EQUATION WITH THE CTS-HS DATA

There are three primary aspects of insurance demand we can examine in this study. The first aspect of insurance demand is the degree of "restrictiveness" of the insurance plan chosen by insured persons. The second aspect of insurance demand is the amount of cost-sharing of the insurance plan chosen. The samples used for these first two aspects of insurance demand are all adults insured in the individual market and all adults insured in the group market. The third aspect of insurance demand which we examine is whether any coverage is obtained. For this analysis, we use samples of the "potential" individual and group markets; the composition of these samples is described further below.

We use pooled data from the 1996-1997 and 1998-1999 Community Tracking Study's Household Surveys (CTS-HS). This nationally representative survey contains health insurance coverage, employment status, and demographic characteristics for a relatively large sample of over 60,000 individuals per period.[1] We choose the CTS-HS since it is both the largest and most recent sample of purchases in the individual market with comprehensive data on the plans obtained. Insured survey respondents answered several questions regarding their plans' restrictions, and a so-called "Followback" Survey of most of the household's insurers was conducted for both confirmation of the reported insurance characteristics and even more detail about the insurance plan. All of our analyses use the sample weights provided with the data so that our results are nationally representative.

Dependent Variables

We examine three different aspects of the demand for insurance for adults. The first aspect is the restrictiveness of the actual plan chosen, and we use four different measures for the restrictiveness of the plan. Specifically, we first estimate a cumulative, or "ordered," logit regression for the sample of insured adults, where the options (in order) are fee for service (FFS), preferred provider organization (PPO), point of service (POS), and health maintenance organization (HMO); this plan typology is the one identified in the CTS-HS Followback Survey. A PPO plan generally uses FFS payments to reimburse providers, but insured individuals are allowed only to choose from a list of "preferred" providers. An HMO often uses capitation to reimburse providers, while a POS plan gives HMO enrollees the option to see providers out of the network at increased fee. The notion is that, if a person were at the margin of desiring a more restrictive plan, the next best choice to a FFS plan would be a PPO, a POS plan instead of a PPO, and an HMO instead of a POS. Specifically, the cumulative logit model (for four discrete choices with discrete probabilities, P_{FFS}, P_{PPO}, P_{POS}, and P_{HMO}, for example) estimates a system of (three, in this case) linear equations for the cumulative log odds:

$$\begin{aligned}
\log((P_{HMO} + P_{POS} + P_{PPO})/P_{FFS}) &= A_{HMO,POS,PPO} + B\,X_i \\
\log((P_{HMO} + P_{POS})/(P_{PPO} + P_{FFS}))) &= A_{HMO,POS} + B\,X_i \\
\log(P_{HMO}/(P_{POS} + P_{PPO} + P_{FFS})) &= A_{HMO} + B\,X_i
\end{aligned}$$

[1] The CTS data samples households across 60 different Metropolitan Statistical Areas and collections of rural counties in the United States Although this unique sampling methodology was utilized to enable local-level identification, the 60 markets were selected in a way so that the resulting CTS-HS sample is nationally representative; for more detail, see Kemper et al. (1996).

where X_i is a vector of person-level demographic characteristics and B is the vector of coefficients of interest; the intercepts for each equation are different but the regression parameters for each logit are constrained to be identical, so that the resulting odds for each of the four discrete choices are proportional.

However, plans of the same type (e.g., HMOs) may still differ considerably in the actual restrictions on care; see, for example, Hacker and Marmor (1999). The second measure for the restrictiveness of the plan that we use, therefore, is the number of actual provider restrictions reported by the respondent. There are three restrictions included in the survey: whether the plan has a directory of doctors who will take its patients, whether the plan requires individuals to sign up with a certain primary care physician, and whether a referral is required to see an out-of-network specialist. Since the latter of these restrictions "builds" upon the former, we estimate a model for the overall number of restrictions, by estimating a Poisson regression for the number of restrictions the insured report; this count of the number of restrictions ranges from zero to three. Because self-reports are not always accurate, the third measure for restrictiveness that we examine is the count of the number of restrictions collected directly from plans in the CTS-HS Followback Survey. (We explore the discrepancies between these measures further below.)

We hypothesize that the results seen in these three measures for restrictiveness will be similar in that each describes managed care restrictions insurers place on enrollees in their ability to obtain (unfettered) medical care. The purpose of analyzing three separate (but presumably highly correlated) measures is to determine whether the results we observe are robust to slightly different specifications (of the dependent variable).

The fourth measure of the plan's restrictiveness is whether the plan includes coverage of mental health services. Ideally, we would examine whether other services are covered by the plan (e.g., prescription drugs, chiropractor), but the breadth of data for services covered by the plan in the CTS-HS Followback Survey (besides the above restrictions on seeing primary care and specialist physicians which we do analyze) is limited to mental health services. We know of no other data source that has more detailed characteristics of plans obtained in both the individual and group markets. Here, we estimate a logit model for the odds that mental health benefits are included in the plan.

The second aspect of the demand for insurance that we consider is the overall amount of cost-sharing in the chosen plan. The CTS "Followback" survey collected data for the annual deductible, the coinsurance percentage (if applicable), and the co-payment (if applicable) required for an in-network provider office visit. Typically, FFS and PPO plans are characterized by deductibles and coinsurance that apply to annual expenditures, while POS and HMO plans have no or minimal annual deductibles but instead require nominal co-payments for each visit. To generate a plan's overall amount of cost-sharing, we therefore use expenditure data from the Medical Expenditure Panel Survey (MEPS) to simulate the average expected out-of-pocket expense for each plan. That is, for each privately insured individual in the MEPS we apply each plan's deductible and coinsurance (if applicable) to their total expenses and/or apply that plan's co-payment (if applicable) to the count of the total number of payments that would be made during the year (making the unlikely, but necessary, assumption that the in-network provider office visit's co-payment in the CTS-HS Followback data was the same for other services, such as inpatient stays and ER visits) to generate an estimate of the plan's out-of-pocket payment and then express that amount as a percentage of total expenditures. Here we estimate an OLS model for the percentage of overall cost-sharing that the plan requires. We expect lower levels of cost-sharing in the group plans than in the individual market plans, due to both the tax subsidy and the lower administrative loading in employment-based groups.

The third aspect of the demand for insurance that we consider is whether individuals (not receiving public insurance) obtain any private coverage of any type. Here we estimate a logit regression for the odds of having any insurance for two samples of all adults who *could* obtain such coverage. For the demand for individual coverage, our sample of so-called "potential" buyers is all adults without public insurance who are in households either where no one is employed or where any employed person is self-employed. For the employment-based group demand equation, the sample is all employees and dependents who are in households where someone is a wage-earner.

However, splitting the sample based on family employment status in this manner is somewhat imprecise because we do observe some mismatches in the types of coverage obtained. Virtually all wage-earners are potential candidates for relatively inexpensive tax-subsidized group insurance coverage since they could decide to work for a firm offering coverage (presumably with relatively lower wages). However, we actually find that a small fraction of wage-earners pay for individual insurance. (Some of them are offered group coverage, while the rest are not.) Perhaps matching to jobs offering coverage is imperfect, perhaps some individuals strongly dislike the group coverage offered (if applicable), or perhaps people are simply behaving irrationally. Only about 15 percent of people with a full-time wage-earner in the family obtain no group coverage and; of those, only 24 percent buy individual insurance coverage. Since the goal of our analysis is to examine differences between individual and group coverage obtained and since the total fraction of people in wage-earning families obtaining individual insurance is only about 3.5 percent, we delete these individual insurance purchasers from the potential group market sample.

More importantly, about 40 percent of people in families where no one is a full-time wage-earner still obtain what is labeled in the survey as employment-based group coverage. Unfortunately, the CTS-HS only includes information for one job per person. Since some individuals hold more than one job (and a self-employed person's secondary job might be with a firm offering group insurance), we are likely classifying some families as nonemployed or self-employed when they actually had access to group insurance. Moreover, since the CTS-HS does not include information for prior employment, some these people could have obtained group coverage earlier (e.g., retiree health, COBRA coverage). Some may have coverage from an alternative group mechanism (e.g., a union of self-employed tradesmen). Assuming that respondents are answering correctly, those who obtain group coverage obviously had an opportunity to obtain group coverage selected by and chosen for the group. As noted earlier, since the purpose of our analysis is to delineate between individual and group coverage, we therefore concentrate on analyses excluding this subpopulation of people who end up with group coverage whom are in our "potential" individual market sample. (We comment on how our results are altered if the sample does not exclude those who had group coverage.) An analysis that examines these illogically classified people may be worthwhile, but this would be outside the scope of our paper.

Explanatory Variables

We are primarily interested in individual-level demographic characteristics related to the demand for insurance. These include income, age, gender, family size, income, education, ethnicity, risk aversion, and health status. Since it is possible that there are regional differences in insurance characteristics that are related to the supply of insurance (and perhaps correlated with differences in the underlying demographic characteristics of the populations), we include binary control variables to identify the eleven Census Divisions of the United States and whether the person lives in an urban area; we also include an indicator for the 1996-1997 period versus the 1998-1999 period.

We define each of our individual-level demographic characteristics with a specification that is as simple as possible so that our test for differences across each characteristic is as strong as possible. It would be possible for us to measure each of these variables in a way that allows for nonlinear effects (i.e., including age and age-squared or categorical variables for various levels of education) and this would undoubtedly increase the fit of our model (since it is likely that such nonlinear effects do indeed exist). However, comparing the difference in the effect of a particular characteristic would be rather difficult with a more complex specification with several terms. Since the primary goal of our analysis is to compare the effects between the individual and group samples, we therefore instead focus on simple specifications for each characteristic so that we may subsequently compare the observed *overall* effect for each particular variable (i.e., a single coefficient).

Since we are theoretically interested in the individual demand for insurance, our empirical model is complicated somewhat by presumed joint choice of insurance plan among people in the household. Ideally, we would like to estimate our models for single-person and family households separately, but we are constrained by the limited sample size of persons in the individual market. We therefore pool single-person and family households together in our analysis, although we do include family size as a demographic characteristic of interest.

We express family income as a percentage of the federal poverty line to adjust for family size and composition, and the log of 100 times this value is used in the estimation.[2] Education is specified as the number of years of schooling, and we consider four categories for ethnicity: Caucasian, Asian, African-American, and Hispanic. We consider two measures of risk-taking behavior. The first is whether the respondent reported strongly agreeing or somewhat agreeing with the statement "I'm more likely to take risks than the average person." The second measure is whether the respondent

[2] Since income is correlated with many of the other demographic measures (as well as possibly endogenous to the generosity of health insurance), we examined models that both excluded and included this income measure. The results that we present for these other variables do not change markedly depending on the inclusion of income.

reported currently being a smoker. Health status is identified by whether the survey respondent reported being in either fair or poor health.[3] The potential endogeneity of health status as a function of insurance coverage (i.e., people who are insured are more likely to obtain medical care and thus report being in better health) is ignored.

Insurance premiums vary considerably because of variation in expected medical utilization and local-area medical care costs, but we are instead interested in the "true" price of insurance—the administrative loading. Unfortunately we do not have accurate direct measures of loading for either individual or group insurance. While we do know that the administrative loading for group insurance varies considerably with the size of the group (and that individual insurance can be identified by a group size of one), this does not provide us with any variation in the administrative loading *within* the individual insurance sample. Moreover, there probably is no exogenous source of loading variation in the individual market. In our base regressions we therefore include no measure of loading for individual insurance.

In the group market as well, we have no direct measure of the average loading for each group. We also do not know how the premium the worker pays (explicit premiums and/or foregone wages) compares with the benefits that worker expects to receive. However, in a second set of models, we do use average firm size in the industry of household employment as a "proxy" measure of the "price" for group insurance. Twelve separate industries are identified in the CTS-HS; we determine the average number of employees in that industry from the 1996 County Business Patterns data. The other source of variation in the net loading—the tax subsidy tied to the household's marginal income tax rate—is obviously very highly correlated with household income, so we assume that the effect of income in the group market picks up both this price effect and any effect of income per se on the demand for insurance. For people in the potential individual market, the only potentially exogenous source of variation in net price is the deductibility of premiums for those with self-employment income. Specifically, the deduction for self-employed people was 25 percent of the premium in 1996 and 45 percent of the premium in 1998. We therefore include a binary variable if a worker in the household was self-employed in this second set of regressions for the effect of price.

Descriptive statistics for insured people in the individual and group market samples are shown in Table 1, and descriptive statistics for the "potential" individual and group market samples are shown in Table 2. These mean values indicate that the plans obtained in the individual market are generally less restrictive than those obtained in the employment-based group market. Individual market plans average 1.4 restrictions on care while group plans average 1.8 restrictions on care. FFS plans comprise 20 percent of individual market plans but only 11 percent of group plans, whereas HMOs are 39 percent of the group market but only 31 percent of the individual market.

[3] We also explored transforming the self-reported health measure (excellent, very good, good, fair, or poor) into a univariate measure of expected expense. Following a procedure outlined in Herring and Pauly (2001) using the Medical Expenditure Panel Survey, we construct a ratio of expected medical expense based on age, gender, location, and health status divided by expected medical expense from only age, gender, and location. Since, the results using this ratio are similar to those using a binary variable for fair or poor health status, we present the latter for simplicity.

TABLE 1
Descriptive Statistics: Insured Persons in the Individual and Group Markets

Variable	Full Sample	Individual Market Sample	Group Market Sample
Family employment type			
No family members are employed	0.070	0.164	0.064
Family contains only a self-employed worker	0.040	0.367	0.020
Family contains a wage-earner	0.890	0.469	0.916
Plan type			
Plan is an FFS	0.111	0.200	0.105
Plan is a PPO	0.367	0.431	0.363
Plan is a POS	0.133	0.060	0.137
Plan is an HMO	0.390	0.309	0.394
Self-reported restrictions			
Plan has a directory of doctors	0.822	0.653	0.832
Plan requires certain primary care physicians	0.567	0.390	0.578
Referral required for out-of-network specialist	0.373	0.314	0.377
Total number of self-reported restrictions	1.762	1.357	1.786
Followback restrictions			
Plan has a directory of doctors	0.907	0.841	0.911
Plan requires certain primary care physicians	0.480	0.321	0.490
Referral required for out-of-network specialist	0.370	0.292	0.374
Total number of Followback restrictions	1.756	1.454	1.775
Followback characteristics			
Mental health covered in the plan	0.986	0.973	0.986
Percent of cost-sharing in the plan	8.244	12.23	7.999
Individual characteristics			
Age	39.96	41.96	39.83
Female	0.519	0.528	0.518
Family size	2.304	2.103	2.316
Family income divided by FPL × 100	442.1	465.9	440.7
Education level	13.82	14.05	13.80
Caucasian	0.787	0.838	0.784
Asian	0.039	0.052	0.039
African-American	0.093	0.040	0.096
Hispanic	0.081	0.071	0.082
Higher-than-average risk taker	0.460	0.478	0.459
Smoker	0.227	0.192	0.230
Fair or poor self-reported health	0.081	0.071	0.082
Number of observations	36,503	2,017	34,486

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.

Presumably these differences largely reflect the historical origins of managed care and especially HMOs as a group insurance product, which has only recently been available in some settings to individuals. In contrast, plans in the individual market are slightly less likely to include mental health services (97.3 vs. 98.6), although the vast majority of plans offer mental health coverage (whether in the individual or group market), due in part to many states mandating that mental health be covered. Thus, our ability to ultimately evaluate the effect of demographic characteristics on the demand for mental health coverage is limited by the low amount of variation among plans.

The amount of cost-sharing is considerably higher in the individual market. About 12 percent of health spending is made out-of-pocket in the individual market, compared to about 8 percent of all spending in the group market. As shown in Table 2, the proportion of people with insurance is considerably higher in the group market relative to the individual market. The lower cost-sharing and higher proportion insured in the group market is surely due to both the tax subsidy and lower administrative loading.

The demographic characteristics are not very different between the sample of insured people in the individual and group markets (as shown in Table 1). In contrast, however, the demographic characteristics are somewhat different between the sample of "potentially insured" people in the individual and group markets (as shown in Table 2) for a few characteristics. Income levels are considerably lower in the potential individual market sample, since they contain the unemployed workers. People in the potential individual market sample also appear to be somewhat more likely to be in poorer health and are more likely to be minorities. However, most of the other demographic characteristics are not very different between samples.

Since the great bulk of people who obtain private health insurance do so in an employment-based group rather than individually, the group market sample size is many times larger than that of the individual market sample. Sample size alone can affect the measured significance level of regression coefficients as well as adjusted measures of goodness-of-fit. Indeed, the significance of the coefficients estimated for our large group sample (presented below) decreases considerably when we estimate models restricting the sample size to that of the individual market sample. Although we do show standard significance levels of 1, 5, and 10 percent, we are therefore somewhat lenient in the significance threshold of the individual market results and the difference in effects between the individual and group samples by discussing results that are significant at 10 percent.

MAIN RESULTS

The results for the three aspects of the demand for insurance that we consider are shown in Tables 3 through 8. Tables 3, 4, and 5 show the results for the "managed care restrictiveness" of the plan chosen by those who are insured. Table 6 shows the results for whether the plan covers mental health services for those who are insured. Table 7 shows the results for the overall amount of cost-sharing for those who are insured. Table 8 shows the results for whether any insurance is obtained. In each of these six cases, a Chow (1976) Test (in results not shown) indicates that the structure of the coefficients for the individual and group market components is significantly different.[4] For this reason we present each measure's results for the two samples of

[4] The Chow Test for the cumulative logit model is performed as a simple OLS model using the values 1, 2, 3, and 4 (corresponding to FFS, PPO, POS, and HMO) as the dependent variable. Similarly, the Chow Test for the Poisson model's count of the number of restrictions uses an OLS model with values of 0 through 3 for values of the dependent variable, and the Chow Test for either any mental health coverage or any insurance coverage is performed as a linear probability model.

TABLE 2
Descriptive Statistics: "Potential" Samples for Individual and Group Insurance

Variable	Full Sample	Individual Market Sample	Group Market Sample
Family employment type			
No family members are employed	0.067	0.590	0.000
Family contains only a self-employed worker	0.046	0.410	0.000
Family contains a wage-earner	0.887	0.000	1.000
Insurance status			
Any private insurance	0.787	0.309	0.848
Individual insurance	0.035	0.309	0.000
Group insurance	0.752	0.000	0.848
Individual characteristics			
Age	38.36	39.73	38.19
Female	0.501	0.521	0.499
Family size	2.230	1.769	2.289
Family income divided by FPL × 100	387.8	275.8	402.1
Education level	13.34	12.41	13.45
Caucasian	0.737	0.636	0.750
Asian	0.107	0.145	0.102
African-American	0.040	0.051	0.038
Hispanic	0.116	0.168	0.110
Higher-than-average risk taker	0.472	0.504	0.468
Smoker	0.271	0.327	0.264
Fair or poor self-reported health	0.102	0.192	0.090
Number of observations	63,550	6,287	57,263

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all adults ages 18 to 64 with no public insurance. People with individual insurance are deleted from the group sample, and people with group insurance are deleted from the individual market sample.

individual and group insurance separately. The final column in each of these six tables shows the statistical significance of the one-tail comparison between each estimated coefficients from the two samples.

Plan Type and Restrictiveness

Table 3 shows results for the cumulative logit model for the plan type; the four plan choices are ordered so that higher parameter estimates from the regression indicate a preference for HMO-type plans over FFS-type plans. Table 4 shows the Poisson regression results for the self-reported number of plan restrictions (with dependent values equaling 0, 1, 2, or 3), and Table 5 shows the Poisson regression results for the Followback number of plan restrictions (0, 1, 2, or 3). Generally, the signs and magnitudes of coefficients are consistent across the three models; more restrictive plans (whether measured by the plan type or the number of restrictions) are (in the combined individual and group market samples) associated with younger age, female gender, smaller families, lower income, minorities, and smokers. Our results for minority enrollment in managed care plans is consistent with that for the "Medicare+Choice" program in

TABLE 3
Demand for Insurance Results: Individual Versus Group Market Comparisons Cumulative Logit for Plan Type Selected by the Insured

Variable	Individual Market Odds Ratio	Group Market Odds Ratio	P-value for One-tail Comparison
Age	1.002	0.992***	0.006***
Female	0.957	1.064***	0.119
Family size	1.024	0.985	0.208
Log family income over FPL x 100	0.953	0.963**	0.410
Education level	0.994	1.006	0.266
Caucasian	n/a	n/a	n/a
Asian	1.240	1.180***	0.403
African-American	1.738**	1.643***	0.402
Hispanic	3.310***	1.735***	0.000***
Higher-than-average risk taker	0.921	0.967	0.294
Smoker	1.323**	1.055**	0.023**
Fair or poor self-reported health	0.986	1.056	0.347
Intercept – HMO	0.216***	0.769**	0.002***
Intercept – HMO or POS	0.290***	1.406***	0.000***
Intercept – HMO, POS, or PPO	2.577**	12.089***	0.000***
Census Divisions included	Yes	Yes	n/a
Maximum-rescaled R-squared	0.169	0.102	
Number of observations	2,017	34,486	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01, and 0.05 (**); between 0.05 and 0.10 (*).

which seniors are given the option whether to enroll in an HMO plan; Thorpe, Atherly, and Howell (2002) find that—when given a choice of whether to enroll—African-Americans are marginally more likely than Caucasians to enroll in a Medicare+Choice HMO (relative to remaining in the traditional fee-for-service program), while Hispanics are much more likely than Caucasians to enroll in a Medicare+Choice HMO.

The most interesting *difference* in coefficients for restrictiveness between the individual and group market samples is with the ethnicity variables. In the "benchmark" individual regressions for plan choice and number of restrictions, Asians and (non-Hispanic) Caucasians who obtain insurance generally have the same demand for coverage restrictions, while Hispanics and African-Americans (holding other variables constant) prefer plans with more restrictions. Hispanics strongly prefer HMO-type plans to FFS-type plans and have a higher number of restrictions in their chosen plans, whether measured by the self-reports or Followback data. African-Americans have a preference for managed care plans but the relationship is not as strong as that for Hispanics. In the group demand regressions for plan choice and number of restrictions, however, minorities are roughly treated the same: Hispanics consistently appear to receive less restrictive plans in the group market than they prefer. There is some evidence that Asians tend to get restricted plans even though they would

TABLE 4
Demand for Insurance Results: Individual Versus Group Market Comparisons Poisson Model for the Insured's Self-Reported Number of Plan Restrictions

Variable	Individual Market Coefficient	Group Market Coefficient	*P*-value for One-tail Comparison
Age	−0.005***	−0.003***	0.198
Female	−0.045	0.017**	0.059*
Family size	−0.030	−0.011**	0.189
Log family income over FPL x 100	−0.017	−0.032***	0.214
Education level	−0.010	−0.004**	0.228
Caucasian	n/a	n/a	n/a
Asian	−0.033	0.110***	0.046**
African-American	0.104	0.118***	0.445
Hispanic	0.253***	0.114***	0.019**
Higher-than-average risk taker	0.027	−0.017**	0.136
Smoker	0.058	0.033***	0.304
Fair or poor self-reported health	0.093	0.033**	0.215
Intercept	0.349*	0.763***	0.015**
Census Divisions included	Yes	Yes	n/a
Pearson Chi-Square/DF	0.961	0.584	
Number of observations	2,017	34,486	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

apparently have been willing to pay to avoid them. Interestingly, the intercept terms are consistently significantly different in the three models, suggesting that all people prefer less-restrictive plans than they actually obtain in the group market.

Table 6 shows results for the logit model for whether the plan obtained includes any coverage for mental health services. Although there is limited variation in whether plans offer mental health services, as 98.6 percent plans offer this benefit, there are a few notable differences between the insurance demand in the individual and group markets. One difference again is for ethnicity. African-Americans and Hispanics appear to prefer plans that do not cover mental health services, but are not able to obtain plans not including mental health services in the group market. Similarly, being older, a smoker, or in fair/poor health implies a weaker underlying preference for mental health services that is not met in the group market.

The Amount of Cost-Sharing

Table 7 shows the results for the OLS model for the overall amount of cost-sharing in the plan obtained. This amount was determined by simulating the total out-of-pocket costs depending on the plan's combination of deductibles, coinsurance, and co-payments. Overall, the signs of the significant coefficients work in the same way in both samples: less cost-sharing is preferred by younger people, African-Americans and Hispanics, and smokers. However, there is again a significant difference in the

TABLE 5
Demand for Insurance Results: Individual Versus Group Market Comparisons Poisson Model for the Insured's "Followback" Number of Plan Restrictions

Variable	Individual Market Coefficient	Group Market Coefficient	P-value for One-tail Comparison
Age	0.001	−0.002***	0.027**
Female	−0.039	0.019**	0.067*
Family size	0.004	−0.008*	0.282
Log family income over FPL x 100	−0.022	−0.013**	0.301
Education level	−0.002	0.003	0.288
Caucasian	n/a	n/a	n/a
Asian	0.063	0.052**	0.446
African-American	0.194**	0.137***	0.267
Hispanic	0.279***	0.121***	0.008***
Higher-than-average risk taker	−0.006	−0.008	0.476
Smoker	0.066	0.021**	0.177
Fair or poor self-reported health	0.070	0.014	0.220
Intercept	0.133	0.520***	0.018**
Census Divisions included	yes	Yes	n/a
Pearson Chi-Square/DF	0.710	0.540	
Number of observations	2,017	34,486	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

coefficients for Hispanics. While the plans Hispanics obtain in the group market do have less cost-sharing relative to non-Hispanic Caucasians, they do not appear to have as little cost-sharing as they would prefer. There appears to be a similar effect for smokers not receiving as little cost-sharing as they appear to prefer, although the low cost-sharing in the individual market may be attributable to higher premiums for smokers rather than to insurance preferences.

Likelihood of Any Coverage

Table 8 shows the odds ratios from the separate logit models estimated for the potential individual and group market samples. Recall that we consider a sample specification for the potential individual market in which our sample excludes those who actually obtained employment-based group insurance. Almost all of the explanatory variables are significantly related to the likelihood of obtaining insurance; those who obtain insurance tend to be older, female, in larger families, higher income, more educated, Caucasian, not risk-takers, nonsmokers, and healthier. (These relationships are consistent with the existing literature on characteristics of the insured versus the uninsured.) The intercept is significantly different, indicating (as would be expected given typical net loadings) that people in the group market, other characteristics held constant, were much more likely to be insured than those who use the individual market.

TABLE 6
Demand for Insurance Results: Individual Versus Group Market Comparisons Logit Model for Whether the Plan Covers Mental Health

Variable	Individual Market Odds Ratio	Group Market Odds Ratio	*P*-value for One-tail Comparison
Age	0.976*	0.999	0.046**
Female	0.907	0.963	0.431
Family size	1.256	0.985	0.118
Log family income over FPL x 100	1.362**	1.006	0.018**
Education level	1.088	1.013	0.157
Caucasian	n/a	n/a	n/a
Asian	0.503	1.243	0.117
African-American	0.364**	1.279	0.010***
Hispanic	0.216***	2.037***	0.000***
Higher-than-average risk taker	0.770	1.071	0.161
Smoker	0.409***	0.839	0.022**
Fair or poor self-reported health	0.291***	1.618**	0.000***
Intercept	999.9	54.01***	0.494
Census Divisions included	yes	yes	n/a
Max.-rescaled R-squared	0.269	0.035	
Number of observations	1,847	32,379	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

For each of these variables, the sign of the effect is consistent between the individual and group market samples. In terms of the odds ratios, all variables with odds ratios less than unity in the group market also have odds ratios less than unity in the individual market, and similarly for odds ratios greater than unity. Given the very large sample sizes, it is not surprising that eight of the twelve coefficients are statistically significantly different (although three of these differences are only significant at 10 percent). However, most of these differences are quantitatively small; for example, the difference in odds ratios for age is only ten percentage points.

However, there are three large differences in magnitude. The difference in the coefficient on family income is easy to understand; the progressive character of the tax deductibility of group premiums exacerbates income difference in the relative likelihood of having insurance. Another main difference is (once again) in the ethnicity variables. Here African-Americans and Hispanics are less likely than Caucasians to have insurance in either setting, but the difference is much more pronounced in the individual than in the group setting.

The third significant difference in the magnitude of the odds of obtaining insurance is related to smoking and being in poor health status. While the direction of the effect is the same in the individual and group markets, the effect of either smoking or being in fair/poor health on being less likely to be insured is stronger in the individual market.

TABLE 7
Demand for Insurance Results: Individual Versus Group Market Comparisons OLS Model for the Percentage of Overall Cost-Sharing in the Plan

Variable	Individual Market Coefficient	Group Market Coefficient	*P*-value for One-tail Comparison
Age	0.039*	0.036***	0.453
Female	0.006	−0.076	0.437
Family size	−0.201	0.041	0.189
Log family income over FPL x 100	−0.244	0.106	0.090*
Education level	0.013	−0.037	0.325
Caucasian	n/a	n/a	n/a
Asian	0.719	−0.226	0.209
African-American	−2.631**	−1.584***	0.212
Hispanic	−3.295***	−1.327***	0.027**
Higher-than-average risk taker	0.090	0.142	0.461
Smoker	−2.003***	−0.235*	0.003***
Fair or poor self-reported health	−1.558	−0.336*	0.113
Intercept	18.70***	8.284***	0.000***
Census Divisions included	Yes	Yes	n/a
Adjusted R-squared	0.154	0.088	
Number of observations	2,017	34,486	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

Since the expectation is that higher expected expenses would either have no effect on the likelihood of obtaining coverage (if premiums are tailored to risk) or have a positive effect on the likelihood of obtaining insurance (if premiums are "pooled" somewhat), the likely explanation for the results we observe is that medical underwriting exists in the both the individual and group markets but that the extent of underwriting is higher in the individual market. As we noted above, using self-reported health status could be somewhat problematic due to its endogeneity to holding health insurance coverage, so these finding are not wholly conclusive. The result that higher expected expenses has a negative effect on group insurance coverage is consistent with our earlier finding that low-income high-risk workers in the small group market are relatively less likely be insured (Pauly and Herring, 1999).

When we used the alternative less plausible measure of potential buyers in the individual market (i.e., adding those who obtained group insurance even though the household appears to only have either nonemployed or self-employed family members), the main results are similar. These results were presented in an earlier version of this paper (Pauly and Herring, 2002) and are available upon request.

Effect of Price on the Demand for Insurance

We re-estimated these regressions adding proxy measures for insurance loading: average firm size in the worker's industry in the case of the group sample and a binary

TABLE 8
Demand for Insurance Results: Individual Versus Group Market Comparisons Logit Model for Any Coverage

Variable	Individual Market Odds Ratio	Group Sample Odds Ratio	*P*-value for One-tail Comparison
Age	1.042***	1.038***	0.092*
Female	1.112	1.141***	0.363
Family size	1.468***	1.573***	0.052*
Log family income over FPL × 100	1.590***	2.360***	0.000***
Education level	1.254***	1.236***	0.183
Caucasian	n/a	n/a	n/a
Asian	0.759*	0.755***	0.487
African-American	0.348***	0.816***	0.000***
Hispanic	0.409***	0.566***	0.005***
Higher-than-average risk taker	0.844**	0.829***	0.404
Smoker	0.435***	0.576***	0.000***
Fair or poor self-reported health	0.542***	0.741***	0.002***
Intercept	0.000***	0.001***	0.080*
Census Divisions included	Yes	Yes	n/a
Max.-rescaled R-squared	0.403	0.323	
Number of observations	6,287	57,263	

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

variable for being self-employed in the case of the individual sample. In results not shown, there was virtually no change in the coefficients or the significance levels for the demand variables we have already discussed.[5] Thus there does not seem to be any correlation between loading proxies and these demand variables and no interactions of note.

The loading proxies themselves were sometimes statistically significant. In the "restrictiveness" regressions, being self-employed (and therefore facing a lower "price" of insurance due to the potential tax deductibility) did lead to a decrease in the restrictiveness of the plan chosen, other things equal. This is what one would expect since deductibility reduced the net price of joining a less restrictive but more costly plan. Being self-employed had no effect on mental health coverage in the individual market, but did (contrary to expectations) increase the amount of cost-sharing, although the effect was small in magnitude. In the group regressions for those who were insured, firm size either had either no significant effect or a very small effect on restrictiveness, mental health, or cost-sharing. The two loading proxies did have large and significant effects on the odds of obtaining any coverage. Other things (including family income) held constant, people were more likely to have insurance if they were self-employed

[5] These results are available upon request.

or worked in industries characterized by larger firms (i.e., where loadings are relatively smaller). Given the different way in which these proxies are specified, we are unable to make comparisons of the magnitude between the individual and group samples.

Goodness of Fit

Overall, the goodness-of-fit is generally not dramatically different (and is relatively low in the restrictiveness models) across the two samples. For the cumulative logit model for plan type (and the standard logit model for both the odds of having mental health benefits and obtaining any insurance discussed below), we use the "maximum-rescaled R-squared" coefficient of determination specified by Nagelkerke (1991) for maximum likelihood estimation models. For the cumulative logit model for plan type, the maximum-rescaled R-squared statistic is 0.169 for the individual market sample and 0.102 for the group market sample.[6] For the Poisson count models for both the self-reported number of plan restrictions and the Followback number of restriction, we use the Pearson's Chi-Square statistic divided by the degrees of freedom to evaluate the fit of the model; values closer to one imply a better fit of the data. For the model of the self-reported number of restrictions, this measure equals 0.961 for the individual market sample and 0.584 for the group market sample. For the model of the Followback number of restrictions, this measure equals 0.710 for the individual market sample and 0.540 for the group market sample. For odds of having mental health benefits, the maximum-rescaled R-squared was 0.269 for the individual market sample and 0.035 for the group market sample.[7] For the amount of cost-sharing, the R-squared is 0.154 in the individual market sample and 0.088 for the group market sample. Finally, for the logit model for the odds of obtaining any insurance, the maximum-rescaled R-squared statistic is 0.403 for the individual market sample and 0.323 for the group market sample.[8] Although the goodness-of-fit is relatively close for each of the individual versus group comparisons for these four models, the goodness-of-fit

[6] Another comparison of the goodness-of-fit for a logit model is to make pair-wise comparisons across observations of the predicted probabilities from the model. Specifically, the ordered logit model produces four predicted probabilities for each individual; the probabilities represent the probability of selecting each of four plans, with the sum of these values equaling one. There are four sets of observations with different values of the observed plan choice. For all pairs of observations with different values of the observed choice, a pair is termed "concordant"if the observed choice has a higher predicted probability than the predicted probability of an observation with a different observed plan choice. The pair is "discordant" if the observed plan choice has a lower predicted probability than that of an observation with a different observed plan choice. A tie occurs if the predicted probabilities for the pair observations are equal. For the ordered logit model, the pair-wise comparisons were concordant or tied with the observed responses for 66.9 percent of the individual market sample and 62.6 percent of the group market sample.

[7] The percent concordant or tied for the predicted and observed probabilities for mental health coverage is 85.4 percent for the individual market sample and 68.2 percent for the group sample.

[8] The percent concordant or tied for the predicted and observed probabilities for having any coverage is 83.7 percent for the individual market sample and 84.8 percent for the group sample.

measures are generally lower for the employment-based group samples, suggesting that there may be unmeasured influences that have different effects in the group setting.

Perceptions of Restrictiveness

For the three "managed care restrictiveness" measures, most of the significant coefficients are similar. However, we thought that any inconsistent effects between the self-reported and "Followback" data were worth exploring. One possibility is that less-educated people mistakenly *think* that there are more or less restrictions than there really are. Or perhaps they prefer more restrictive plans but are unable to accurately select them. Another possibility was that "Followback" Survey of insurers (conducted a year later than the Household Survey) may have introduced some errors into the restrictions variables. To examine this more closely, we estimated a Poisson regression for the number of discrepancies between the self-reported restrictions and the Followback restrictions; with three restrictions, the number of discrepancies ranges from 0 to 3. As shown in Table 9, income and education are both significantly negative, supporting the notion of confusion among the low-income/low-educated. However, males, those in smaller families, African-Americans, Hispanics, risk-takers, and smokers (controlling for income and education) were also more likely to have discrepancies. However, whether the plan was obtained in the individual market and the regional binaries for the nine Census Divisions were also significant. This suggests that the Followback Survey of insurers (conducted a year later than the Household Survey) may have introduced errors into the restrictions variables, as fewer discrepancies were found in New England and the Mid-Atlantic, which is where the administrators of the survey are located. Nevertheless, the overall conclusion is that, whatever the effects of income and education on the demand for plan restrictions (or on whether people actually obtain the plan they believe they hold), these effects do not differ by insurance choice setting.

THE EFFECT OF UNIONIZATION ON GROUP DEMAND

Here we examine whether unionization is associated with the level of coverage obtained for those in the group market. We therefore reexamine the six different measures of the demand for insurance discussed above, but now include an exogenous measure of unionization in the regression. Although the CTS-HS data does not indicate union status for workers, whether a worker is actually a member of a union would be endogenous in any case, as noted above. We therefore obtain a measure of the local-level prevalence of unions in the worker's particular industry. (We assume that individuals choose their job-specific skill level—applicable to a particular industry—independent of their preference for health insurance.) Accordingly, we now limit our relevant samples to workers only; the sample above included nonworkers. Twelve separate industries are identified in the CTS-HS. The local-level measure of unionization for the worker's particular industry is obtained from both the 1996 and 1998 Current Population Study (CPS). We determine the proportion of workers in unions for the twelve industries identified in the CTS-HS for each of the nine Census Divisions

TABLE 9
Poisson Model for the Number of Discrepancies Between the Self-Reported and "Followback" Plan Restrictions

Variable	Coefficient
Individual insurance plan	0.055**
Age	−0.001
Female	−0.024*
Family size	−0.020***
Log family income over FPL x 100	−0.034***
Education level	−0.022***
Caucasian	n/a
Asian	0.020
African-American	0.074***
Hispanic	0.049**
Higher-than-average risk taker	0.022*
Smoker	0.032**
Fair or poor self-reported health	−0.020
Year 1998	0.008
Metro	0.007
New England	n/a
Mid-Atlantic	0.029
East North Central	0.145***
West North Central	0.076
South Atlantic	0.104***
East South Central	0.156***
West South Central	0.082***
Mountain	0.098***
Pacific	0.115***
Intercept	0.124*
Pearson Chi-Square/DF	0.795
Number of observations	36,503

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured adults ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

in the United States during each of these years.[9] This local-level unionization for one's industry is the exogenous measure (varying both across markets and over time) of union status we use for our analysis. The mean for this unionization measure across all individuals is 0.155, and its standard deviation is 0.159. This is, by construction, consistent with approximately 15 percent of workers in the United States being members of unions.

[9] Since only a fraction of workers in the CPS have data for union status (as a part of the earnings questionnaire), sample size considerations prohibit us from determining MSA-level measures of unionization by industry.

TABLE 10
Effect of Unionization on Workers' Demand for Insurance

Variable	Individual Market Result for Unionization	Group Market Result for Unionization	*P*-value for One-tail Comparison
Ordered Logit's odds ratio for plan type selected by the insured	2.067*	0.848**	0.013**
Poisson model's coefficient for the insured's self-reported number of plan restrictions	0.049	−0.045*	0.291
Poisson model's coefficient for the insured's "followback" number of plan restrictions	0.152	−0.061**	0.103
Logit model's odds for whether the plan covers mental health	131.6*	3.680*	0.080*
OLS model's coefficient for the percentage of overall cost-sharing in the plan	−1.328	1.095***	0.164
Logit model's odds ratio for having any coverage	0.943	5.730***	0.000***

Source: 1996-1997, 1998-1999 Community Tracking Study Household Surveys.
Notes: Samples include all insured workers ages 18 to 64 with no public insurance.
P-values: Statistical significance at 0.01 or better (***); between 0.01 and 0.05 (**); between 0.05 and 0.10 (*).

It may still be the case that people who work in unionized industries and/or regions have special tastes for health insurance that are unobservable (i.e., not measured by our set of individual-level demographics). For this reason, we also estimate the effect of our unionization measure on the demand for insurance in the individual market sample. A conservative test for the impact that unions have on affecting *group* choice is then measured as the *difference* of the coefficients for the group market sample relative to the individual market sample.

Unionization Results

Table 10 shows the demand for insurance results (using the six different dependent variables) for the two samples as a function of the local-level unionization for one's industry. The results indicate that unionization is associated with increased plan generosity for those in the group market in two main ways. First, unionization is associated with fewer managed care restrictions to be obtained by those obtaining group insurance; this negative effect is seen for plan choice and both the self-reported and Followback restrictions. For plan choice, the difference between the individual and group markets is significant. For the number of restrictions, even though the unionization measure is insignificant for the individual market samples, the difference in these two coefficients (between the individual and group market samples) is not statistically significant, so we cannot rule out unmeasured tastes correlated with unionization rather than unionization per se. However, it is probable that the insignificant comparisons are most likely driven simply by the large standard errors within the small individual market sample's estimation.

Second, unionization is strongly associated with an increase in the likelihood of obtaining any coverage for those in the group market. The coefficient for this measure in the group model is significantly positive *and* the magnitude of the group sample's coefficient is significantly *larger* than that for the individual market sample. While the unionization measure was marginally significantly positive for the likelihood of covering mental health services in the group market, there is a huge coefficient in the individual market which is unexplainable (and perhaps simply an aberration resulting from the small amount of variation in mental health benefits observed for such a small sample). Unionization also is positively associated (contrary to expectations) with the amount of cost-sharing in the group market, but there is not a significant difference between the individual and group markets for cost-sharing. Unionization did not affect the signs or magnitudes of the other influences on demand in any of the regressions, and the patterns we observe above are generally the same in this second set of results in which the sample is limited to workers.

DISCUSSION

Many of the significant and obvious influences on insurance choice have what we judge to be similar effects in the group setting as in the "benchmark" individual choice setting. Although the statistical tests suggest, and inspection of individual coefficients confirms, that there are some differences, we conclude informally that the overall pattern is fairly similar based on two findings. First, the direction of effects is overwhelmingly the same in both individual and group specifications (with only ethnicity standing out as an exception). Age, female gender, income (except when affected by the tax subsidy), and education all influence demand in the same direction. Second, many coefficients are similar in magnitude and often not statistically significantly different. Where there are differences that are significant statistically, it is not clear whether they are different to a great enough extent to matter for policy since we are not sure what the policy concern would be. For example, would disparities related to age on the number of plan restrictions be of great concern? One important caveat of our study, however, is that we are unable to examine all aspects of insurance demand (e.g., coverage of prescription drugs and chiropractors) due to a lack of data availability, so we cannot be sure that differences do not exist for other aspects of insurance demand.

These results definitely reject the "random assignment" model of insurance coverage and confirm the hypothesis that individuals with different preferences tend to sort into firms with different benefit structures. The economic variables seem to affect demand approximately equally; matching on these variables appears to be fairly good. But whether the choices are close enough to mollify those critics of group coverage who attach high value to insurance that individuals choose and "own" is not clear.

The main difference attributable to the group setting that does have policy relevance has to do with ethnicity and unionization. In group settings, Hispanics obtain plans that are less restrictive, are more likely to cover mental health services, and have higher amounts of cost-sharing than they generally prefer. African-Americans also obtain plans in the group setting that are relatively more likely to cover mental health services than they prefer. Hispanics and African-Americans are both likely to have insurance more frequently than they prefer. Unionization also appreciably affects insurance, leading to more coverage and less restrictive benefits. These findings for unionization are consistent with the theoretical hypothesis suggested by Goldstein and Pauly (1978). There is some evidence that this "taste" spills over into nonunionized workers in heavily unionized industries.

Considering the relative complexity of the task of linking group benefits to individual preferences, our judgment is that a reasonable amount of matching does occur. People who want generous, costly plans tend to get them. People who choose more restrictive plans or no insurance at all when they have an individual choice also appear to select jobs with the same limits, presumably offset by higher wages. There still surely can be some individual workers trapped into mismatches: either higher wages but less generous benefits than they want, or higher benefits but lower wages in other settings. There are some cases of imperfection, especially as related to ethnicity and unionization. The next step would be to investigate what firm or individual characteristics are associated with mismatches. The "job-lock" literature does this for the individual characteristic of poor health, and further investigation of this phenomenon seems warranted.

REFERENCES

Blumberg, L., and L. Nichols, 2004, Why Are So Many Americans Uninsured? A Conceptual Framework, in *Health Policy and the Uninsured* (Washington, DC: Urban Institute Press).

Blumberg, L., and M. Cancian, 2004, Do Workers Have a Choice and Sort Strategically? Evidence from an Analysis of Spouses' ESI Offers, Urban Institute Working Paper.

Browne, M., 1992, Evidence of Adverse Selection in the Individual Health Insurance Market, *Journal of Risk and Insurance*, 59.1: 13-33.

Bundorf, M., 2002, Employee Demand for Health Insurance and Employer Health Plan Choices, *Journal of Health Economics*, 21: 65-88.

Cardon, J., and I. Hendel, 2001, Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey, *RAND Journal of Economics*, 32.3: 408-27.

Chow, G. C., 1976, Tests of Equality between Sets of Coefficients in Two Linear Regressions, *Econometrica*, 28: 591-605.

Goldstein, G., and M. Pauly, 1976, Group Health Insurance as a Local Public Good, in: R. Rosett, ed., *The Role of Health Insurance in the Health Services Sector* (New York: National Bureau of Economic Research).

Hacker, J., and T. Marmor, 1999, The Misleading Language of Managed Care, *Journal of Health Politics, Policy and Law*, 24: 1033-1044.

Herring, B., and M. Pauly, 2001, Premium Variation in the Individual Health Insurance Market, *International Journal of Health Care Finance and Economics*, 1: 43-58.

Kemper, P., D. Blumenthal, J. Corrigan, P. Cunningham, S. Felt, J. Grossman, L. Kohn, C. Metcalfe, R. St. Peter, R. Strouse, and P. Ginsburg, 1996, The Design of the Community Tracking Study: A Longitudinal Study of Health System Change and its Effects on People, *Inquiry*, 33: 195-206.

Monheit, A., and J. Vistnes, 1999, Health Insurance Availability at the Workplace: How Important are Worker Preferences?, *Journal of Human Resources*, 34: 770-785.

Moran, J., M. Chernew, ed., and R. Hirth, 2001, Preference Diversity and the Breadth of Employee Health Insurance Options, *Health Services Research*, 5: 911-934.

Nagelkerke, N., 1991, A Note on a General Definition of the Coefficient of Determination, *Biometrika*, 78: 691-92.

Pauly, M., and B. Herring, 1999, *Pooling Health Insurance Risks* (Washington, DC: AEI Press).

Pauly, M., and B. Herring, 2002, The Demand for Insurance in the Group Setting: Can You Always Get What You Want? Presented at NBER Health Care Program, May 3, 2002.

Rothschild, M., and J. Stiglitz, 1976, Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information, *Quarterly Journal of Economics*, 90.4: 630-649

Thorpe, K., A. Atherly, and K. Howell, 2002, *Medicare+Choice: Who Enrolls?* Paper prepared for the Blue Cross and Blue Shield Association.

Tiebout, C., 1956, A Pure Theory of Local Expenditures, *Journal of Political Economy*, 64: 416-424.