Social Learning in the Demand for Employer-Sponsored Health Insurance

Alan T. Sorensen*

October 5, 2001

Comments Welcome

Abstract

I use data from the University of California to empirically examine the role of social learning in employees' choices of health plans. The basic empirical strategy starts with the observation that if social learning is important, decisions should appear to be correlated across employees within the same department. I present evidence of considerable "clustering" on health plans within departments, and estimate discrete choice models in which individuals' perceptions of the payoffs associated with different alternatives are influenced by peers' decisions. To distinguish social learning from the potential influence of common unobservable characteristics I propose and conduct a variety of tests, which taken together favor the social learning explanation. For example, I find that newly hired employees are more responsive to their peers' decisions than existing employees, which I interpret as reflecting their relative lack of private information (and consequently their relative reliance on information communicated from coworkers).

1 Introduction

Individuals often have incentives to learn from their neighbors before making economic decisions. In some cases, other individuals' choices serve as signals of private information, so learning comes from merely observing their actions—for instance, a typical tourist will (rationally) avoid empty restaurants and prefer those that are crowded with locals. In other cases, information and experiences are shared directly through conversation, as when a consumer planning to buy a car asks her friends about their experiences with different brands or different dealers. A role for social learning exists whenever grouped agents use independent information to make decisions involving uncertain payoffs.

A growing body of theoretical research has incorporated social learning into standard models of economic decision-making. One thread in this literature shows that when individuals make decisions sequen-

*Department of Economics, University of California, San Diego, and NBER; asorensen@ucsd.edu. I thank Michele French and the University of California for providing the data, Zhigang Li for outstanding research assistance, and the National Science Foundation for financial support (grant SES-0079201). Seminar participants at UC Davis and the University of Arizona provided helpful comments. Any errors are mine.

tially, they may rationally choose to ignore their own private information in favor of mimicking their predecessors, resulting in “herd behavior” (Banerjee 1992) or “informational cascades” (Bikhchandani, Hirshleifer, and Welch 1992). Other papers have explored the dynamic implications of social learning for technology adoption (Ellison and Fudenberg 1993) and for the prevalence of conformity vs. diversity in a population (Bernheim 1994; Ellison and Fudenberg 1995).

The objective of this paper is to empirically identify the role of social learning in a specific context: individuals’ choices of employer-sponsored health plans. Using data from five University of California campuses, I first present evidence of “clustering” on health plans within departments: the plan choices of employees in the same department are “too similar” relative to what we’d expect based on individual characteristics and campus-wide patterns. I then estimate an econometric model of health plan choice that explicitly allows for social learning, and find large, statistically significant effects that are robust across campuses and model specifications.

This work follows closely in the spirit of several empirical papers that have examined social learning hypotheses in other settings. Glaeser et. al. (1996) examine the role of social interactions in individuals’ decisions to commit crimes, and estimate indexes of social interaction for different types of criminal activity. Their study is motivated by the high cross-city variance in crime rates (too high to be explained by observable city characteristics), which is analogous to the finding here that departments exhibit too much clustering on health plans to be explained by compositional characteristics of the departments themselves. This paper is also similar to a recent study of Duflo and Saez (2000), who also look at data from a university and exploit departments as reasonable proxies for communication networks. Their study finds significant evidence that individuals’ decisions about whether to enroll in a university-sponsored retirement plan are influenced by the decisions of their peers within the same department. Other empirical studies have addressed the impact of social learning on the adoption of fertilization technologies in agriculture (Conley and Udry 2000; Munshi 2000), labor supply decisions (Woittiez and Kapteyn 1998), welfare program participation (Bertrand, Luttmer, and Mullainathan 2000), stock market participation (Hong, Kubik, and Stein 2001), and the use of contraception (Munshi and Myaux 2000).[1]

[1]Brock and Durlauf (2001) also analyze discrete choice with social interactions, but their approach is somewhat different from the one employed here. They consider games of binary choice in which payoffs depend on neighbors’ choices, but in which there is no direct communication or cooperation (so that individuals make choices based on expectations of their peers’ decisions). The authors use tools from statistical mechanics to analyze self-consistent equilibria, pointing out that there can exist multiple equilibrium average choice levels. In contrast, for employees’ health plan decisions we expect choices to be made after direct communication between peers within a department (so that at least some peers’ choices are known with certainty), and uncertain forecasts of coworkers’ decisions do not play an important role. Nevertheless, an interesting direction for future research would be to extend the Brock and Durlauf approach to polychotomous discrete choice environments so it can be applied to the present context.

All of these papers confront some common identification problems, many of which have been described in detail by Manski (2000). Of principal relevance for this paper is the difficulty distinguishing between what Manski calls "endogenous interactions," in which individual decisions are influenced directly by the decisions of their peers, and "correlated effects," in which the decisions of individuals within a group are similar due to shared (and possibly unobservable) characteristics. Other authors have attempted to resolve this issue in a variety of ways: for example, Munshi and Myaux (2000) show that women's decisions about contraception respond strongly to the decisions of other women within the same religious group, but not to the decisions of women in a separate religious group within the same village. The absence of cross-subgroup effects suggests interpreting the measured own-group effects as social interactions rather than spurious correlation due to common unobservable characteristics. Duflo and Saez (2000) use an instrumental variables strategy to address the endogeneity problem potentially induced by common unobservables. They regress a dummy for participation in a university retirement plan on the average participation within the employee's department, using average department wages as an instrument for average participation.[2]

I attempt to resolve the question of social effects vs. common unobservables using a number of tests and strategies. The first of these exploits the availability of data across different UC campuses. If correlation in behavior is driven by common unobservables, one should expect these unobservables to also be correlated across campuses: for example, if economists at UC San Diego share a common (but unobservable) characteristic that causes their health plan choices to be similar, their choices should also be similar to those of economists at Berkeley, Davis, Los Angeles, and Irvine. Since I have data at several campuses, I can test for the presence of department-specific effects simply by trying to estimate them. The evidence from this test is mixed: in some rare cases department-specific effects appear to be statistically significant, but in the majority of cases these effects are indistinguishable from zero. The second robustness check follows the suggestion of Munshi and Myaux by subdividing academic departments into faculty and staff. Estimated own-group effects are large and statistically significant, while the cross-group effects are in most cases indistinguishable from zero—a result that is clearly consistent with the social learning hypothesis that communication within academic departments tends to be isolated within faculty and staff subgroups. Finally, I consider the behavior of newly hired employees and employees who switch plans, presenting an argument that the former should appear to be more sensitive to peers' choices, and the latter less sensitive; the data are shown to be generally consistent with these hypotheses.

[2] Duflo and Saez also conduct tests similar to the one proposed by Munshi and Myaux, dividing university departments into subgroups based on age, gender, etc., and comparing the estimated own-group effects with the cross-group effects. Like Munshi and Myaux, they find that cross-group effects are generally absent, and take this as evidence in favor of the social effects explanation for the apparent correlation in behavior.

2 Background and Data

The University of California (UC) system is comprised of 9 university campuses plus four additional research laboratories. Nearly all full-time employees and some part-time employees are eligible to enroll in one of the health plans offered through the UC benefits program. The typical employee at a UC campus can choose from one of three HMOs (Health Net, Kaiser, and Pacificare), a point-of-service (POS) plan (UC Care), and a traditional fee-for-service plan (Prudential High option). The HMO plans typically require little or no out-of-pocket payments from the employee, while the POS plan requires monthly out-of-pocket payments ranging from \$17-\$50 (depending on the number of dependents to be covered under the employee's plan). Enrollment in the fee-for-service plan is very rare, since the out-of-pocket payments required are on the order of \$1,000 per month. Employees who choose not to enroll in one of the available plans are automatically given minimal coverage through a default "Core Medical" option. Plan enrollments by campus are shown in table 1.

The data used in this study were provided by the UC benefits office. The data cover employee health plan decisions for the years 1995-2000 at each of the nine university campuses; I will focus attention on the five largest campuses: Berkeley, Davis, Irvine, Los Angeles, and San Diego. For each employee, the data indicate the plan chosen by the employee, the department in which the employee works, the date the employee was hired, and the employee's monthly salary, along with additional demographic characteristics including age, sex, and zip code of residence. Family status can be roughly inferred from the coverage type chosen with the health plan (single-party, two-party, or family). The availability of relatively rich demographic information is critical in this study, in particular since we expect individual-level heterogeneity in preferences over health plans to be driven largely by differences in age, income, and place of residence. Price considerations play a diminished role in health plan decisions: though the decision of whether to enroll in the POS plan or an HMO may be driven by cost concerns (the POS plan costs \$15-\$45 more per month than the HMOs), decisions among the three HMOs are based on non-price considerations since the required employee premiums are essentially the same (zero). The set of available options and the corresponding pricing structure has remained constant over the sample period for most campuses, with only a few minor changes.[3]

Each campus has a benefits office responsible for disseminating information about the available health insurance options. Newly hired employees are encouraged to attend orientations in which the plans (along

[3] A substantial change in the pricing structure occurred just prior to our sample period; see Buchmueller and Feldstein (1997) for an analysis of the impact of the price change on health plan switching using data very similar to the data used here.

with other employee benefits) are explained, and brochures with basic information are typically mailed to employees prior to periods of open enrollment. Most departments in the university system have a staffperson assigned as the benefits coordinator who may serve as the point person for information within departments.

Most of the institutional arrangements for disseminating health plan information focus on informing employees about the process of enrolling or about the plans' relative costs and payment structures (e.g., explaining the difference between a point-of-service plan and an HMO). From an employee's perspective, much of the relevant information needed to choose among plans—e.g., the quality and geographic locations of the plans' doctors, the incidence of reimbursement hassles, the difficulty of getting access to specialists, etc.—cannot (or at least is not) disseminated systematically. As a consequence, much of the information used as an "input" to the health plan decision may be learned through word-of-mouth communication and conversation within peer groups. This social learning aspect of the health plan decision should lead to an observed similarity in the decisions of employees sharing a common social network; in the following section I present basic evidence on the degree of observed similarity within departments, which I take as an appropriate proxy for defining an employee's social network.

3 Evidence of clustering within departments

Although social learning about health plans can take a variety of different forms, this paper will presume that any form of communication among peers will cause their decisions to be positively correlated. The simplest example of this effect occurs when individuals, facing uncertainty about the relative payoffs of the available options, draw inferences about plan qualities by observing the choices of their peers. If the individuals' own private information about the plans is sufficiently weak, this sort of learning can lead to mimicry (herding), which is the most obvious form of correlated behavior that can arise from peer effects. Note, however, that pure mimicry is not necessary for social interactions to generate positively correlated decisions. If employees within a group share common information (e.g., through casual conversations and word-of-mouth communication), their decisions will tend to be similar even if they don't directly observe or mimic each other's actual choices.[4]

In this section I present basic evidence that employees' health plan decisions are correlated within departments. To the extent that departments can be regarded as the relevant social networks in which indi-

[4]While it is possible to tell stories about how some forms of communication could lead to *negatively* correlated decisions, I doubt that such effects could be predominant. For example, if Jane is enrolled in Health Net and expresses dissatisfaction in a conversation with her peers, her peers will be more likely to choose plans different from Jane's (negative correlation). However, all of the peers will be acting on a common "information input," so their decisions will tend to be positively correlated. More importantly, if Jane is dissatisfied with Health Net and expresses that to her peers, she is likely to act on that information herself by switching plans.

viduals communicate about health plans, such evidence should be seen as consistent with a social learning hypothesis. Note, however, that merely looking for correlation in employees' decisions is a low-power test of the social learning hypothesis: the absence of any correlation would allow us to reject the hypothesis, but correlation may be present even if social learning is completely irrelevant. Testing the social learning hypothesis against other plausible alternatives is an issue I take up later in the paper.

To measure the degree of similarity among co-workers' health plan decisions, we can borrow an approach from the literature on industry agglomeration that maps almost perfectly into the present analysis. In the presence of local knowledge spillovers, the payoffs to siting a plant in state or city A increase with the number of other plants also located in A, potentially resulting in agglomeration. To measure the degree of agglomeration in an industry, Ellison and Glaeser (1997) propose starting with

$$G_j = \sum_{k=1}^{K} (s_{jk} - x_k)^2 \ , \tag{1}$$

where k indexes locations, s_{jk} is the share of industry j's employment located in location k, and x_k is the share of total employment (across all industries) located in k. G_j is obviously maximal if all of industry j's employment is agglomerated in one place, and minimal (0) when the pattern of employment in industry j exactly matches average employment patterns across all industries.

The function G can be used in an analogous way to measure the degree of health plan "clustering" within a department: just let k index health plans and j index university departments. If every employee in a department chooses the same health plan, that department will have a high G, while G will be low if the within-department shares closely match the shares over the entire university. Note, however, that taking x_k to represent the share of plan k in the entire university ignores differences in demographic composition across departments. For instance, if employees in the English department are young and poor relative to the average university employee, and if age and income are important determinants of health plan choice, then it would not make sense to expect the health plan shares in the English department to match the plan shares for the university as a whole. To account for the importance of individual characteristics (and their composition within departments), we can simply replace the x_k with an appropriate function of observable characteristics. For department j, we can estimate

$$G_j = \sum_{k=1}^{K} \left(s_{jk} - \frac{1}{n_j} \sum_{i=1}^{n_j} p_k(z_i) \right)^2 = \sum_{k=1}^{K} \left[\frac{1}{n_j} \sum_{i=1}^{n_j} (d_{ik} - p_k(z_i)) \right]^2 \ , \tag{2}$$

where n_j is the number of employees in department j, d_{ik} is an indicator equal to one if employee i chose plan k, and $p_k(z_i)$ represents the probability that an individual with characteristics z_i chooses plan k.

To see that positive correlations in employee decisions will lead to higher indexes of clustering, note that

$$\begin{aligned} E[G_j] &= E\left[\sum_k \left(\frac{1}{n_j^2}\sum_{i=1}^{n_j}(d_{ik}-p_k(z_i))^2 + \frac{2}{n_j^2}\sum_{i=1}^{n_j}\left[(d_{ik}-p_k(z_i))\sum_{l=i+1}^{n_j}(d_{lk}-p_j(z_l))\right]\right)\right] \\ &= \sum_k \left(\frac{1}{n_j^2}\sum_{i=1}^{n_j} E\left[(d_{ik}-p_k(z_i))^2\right] + \frac{2}{n_j^2}\sum_{i,l:i\neq l} E\left[(d_{ik}-p_k(z_i))(d_{lk}-p_j(z_l))\right]\right) \end{aligned} \quad (3)$$

The second term is a covariance term: if d_{ik} and d_{lk} are independent, this term will be zero, whereas if Cov$(d_{ik}, d_{lk}) > 0$, this term will be positive and we expect G to be higher. Here it is important to note that social effects are not the only factors that could induce positive correlation in the decisions d_k. If employees in department j share common characteristics that influence their choices but are *not* accounted for in the characteristics vector z, this will also lead to Cov$(d_{ik}, d_{lk}) > 0$. If we find that measured G indexes are high, this will indicate that behavior among employees in a department is positively correlated, but will not distinguish between social learning and common unobservables as explanations for the correlation.

The critical question for our purposes is whether the degree of observed health plan clustering is excessive relative to what we'd expect in the absence of social effects or common unobservables. Ellison and Glaeser address the analogous question for studying industry agglomeration: are industries more geographically clustered than we would expect them to be if they chose plant locations by "throwing darts at an appropriately scaled map?" To answer the question, they specify a parametric model generating employment shares (s), derive the expectation of G as a function of the parameters, and then propose an index of agglomeration whose expectation is zero in the absence of agglomerative forces (local spillovers or natural advantages) and increases in the strength of these forces.

Since similarity in the decisions of employees in a given department may merely reflect relatively homogeneous individual characteristics in that department, measures of observed plan concentration should be compared to appropriate, department-specific benchmarks that account for the composition of employee characteristics. The approach of Ellison and Glaeser can be readily translated to the current setting, but instead I propose an alternative method that more easily accounts for departmental differences in demographic composition. For each department, consider bootstrapping a test of the hypothesis that the observed measure of clustering (G) was the result of purely independent individual decisions—i.e., the hypothesis that neither social effects nor common unobservables are important. We can simulate the distribution of G under this hypothesis by having each employee's plan choice be a draw from the multinomial distribution, with the probability of choosing plan k given by $p_k(z_i)$ (the probabilities depending only on the employee's own characteristics z_i). The question then arises: what shall we use for the $p(z)$ functions? For any given set of observable characteristics z, in principle the probabilities $p_k(z)$ could be computed as cell averages. However, if z contains a large number of characteristics, one would need an enormous dataset to reasonably estimate $p_k(z)$ using cell averages. Instead, I propose estimating a multinomial logit model in which plan utilities depend on the characteristics z, and substituting the predicted probabilities from this model for the $p_k(z)$ functions. Once the distribution of G has been simulated under the null hypothesis of no social effects, we can simply observe where the observed value of G falls in the distribution, rejecting the null hypothesis if the observed value falls in the upper tail of the bootstrapped distribution.

Results from this procedure (for departments in each of the five largest UC campuses) are reported in table 2. The details of the process are as follows. First, for each campus I separately estimate a multinomial logit model in which an employee's choice of plan is a function of age, salary, tenure (time since hired), sex, zipcode of residence,[5] and family status.[6] Second, for each department at each campus I simulate the distribution of G in the absence of social effects by having employees in the department choose plans independently with probabilities given by $\hat{p}_k(z_i)$, the predicted probabilities from the multinomial logit model, and computing G from the resulting choices. Repeated b times, this procedure yields a sequence $G_j^1, \ldots, G_j^b$ of indexes for department j. Finally, the actual (observed) value of G is computed for each department as in equation 2 (again substituting the predicted probabilities from the multinomial logit model for $p_k(z_i)$), and the bootstrapped p-value for the hypothesis of no social effects or common unobservables is given by $\frac{1}{b}\,\#\left\{G_s : G_s > G_{observed},\ s = 1, \ldots, b\right\}$. For the results reported here, the number of bootstrap repetitions (b) was 10,000.

The bootstrapped p-values of the observed G indexes are summarized in table 2. For each campus, the table lists the fraction of departments whose actual value of G falls in the specified range of the bootstrapped distribution. The results of the simulations indicate that measures of clustering are extremely high for a relatively large fraction of departments. Perhaps the easiest way to interpret these results is to see them as bootstrapped tests of the hypothesis that plan choices are made on a purely individual basis, with the choice probabilities being functions of observable employee characteristics. For departments with p-values less than 0.05, we would reject that hypothesis at the 5% level. If the hypothesis of no correlations in behavior were true, we should expect to reject the hypothesis for roughly 5% of the departments as a consequence

[5]The number of different zipcodes ranges from roughly 70-180, depending on the campus. These are incorporated in the multinomial logit model simply as a full set of zipcode dummies.

[6]These first-stage logit models substantially improve the predictions of plan shares relative to a model based solely on university-wide averages. The pseudo-R^2 measures are .09, .09, .11, .14, and .11 for Berkeley, Davis, Irvine, Los Angeles, and San Diego, respectively.

of sampling error. Instead, we find that at each campus we would reject the hypothesis for a substantially larger fraction of the departments (between 18 and 33 percent, depending on the campus). If we look at the uppermost decile (p-values less than 0.10), between 27 and 43 percent of the departments fall in this extreme. The numbers are roughly similar across campuses, but they are most striking for UCLA: nearly 33 percent of the UCLA departments have G's in the upper 5% tail of the simulated distribution, and nearly 64 percent are in the upper quartile.

Table 3 gives some examples from UC San Diego to illustrate the story behind the numbers of table 2. The Economics and Linguistics departments both have very low G indexes relative to what we'd expect based on the bootstrapped distributions, reflecting the fact that their actual plan shares closely match the predictions from the multinomial logit model with independent, individual characteristics-driven choices. Conversely, the health plan choices in the Physics and Cellular & Molecular Medicine departments are significantly more clustered than we would expect. The physicists are gravitating too much toward UC Care to be explainable by individual demographics and sampling error; employees in the Cellular & Molecular Medicine department are clustering too heavily on Pacificare. It is also evident from the table that the multinomial logit model generates predictions that can vary substantially across departments (note, for example, the differences in predicted plan shares for UC Care), suggesting that it would be a mistake to simply use campus-wide plan shares as proxies for expected plan shares within departments.

Figures 1 through 5 provide a visual illustration of the prevalence of clustering at each campus. If social effects and common unobservables are unimportant, the bootstrapped percentile of each department's G (i.e., its relative position in the bootstrapped distribution of G) should be a random draw from the $U[0,1]$ distribution. The empirical distribution of these percentiles should therefore look like a horizontal line at 1. The figures show histograms of the department percentiles for each campus, along with the uniform density superimposed. The skewness of the empirical distributions indicates that observed levels of clustering exceed expectations too often to have been generated by the aggregation of purely independent individual decisions.

Figure 6 demonstrates the importance of accounting for departments' demographic compositions in assessing the degree of health plan clustering. The top panel shows the distribution of G-percentiles at UC San Diego when plan shares are predicted using employee characteristics (this is just figure 5 drawn on a different scale); the bottom panel shows the distribution that would result if predicted plan shares were simply the campus-wide plan shares. The "uncorrected" distribution is much more severely skewed: failing to account for employee characteristics would lead us to overstate the degree of health plan clustering within departments, since apparently much of the similarity of health plan choices within departments arises from similarity in the employees' characteristics.

4 Econometric Models and Results

The evidence presented in section 3 indicates that individuals' decisions are correlated within departments. In order to measure the impact of co-workers' choices on individuals' choices, in this section I specify and estimate discrete choice models in which peer effects are explicitly incorporated. As a starting point, consider a model in which employee i in department j gets a random payoff from choosing plan k equal to

$$u_{ijk} = x_i' \beta_k + f(d_{-i,jk}) + \eta_{jk} + \epsilon_{ik} \ , \tag{4}$$

where x_i is a vector of observable characteristics (such as age, income, etc.), η_{jk} is an unobservable preference for plan k that is shared by all employees in department j, and ϵ_{ik} is a random preference shock representing the idiosyncratic taste of employee i for plan k. The potential influence of peers' decisions is captured in the function $f(d_{-i,jk})$, where the argument is the vector of decision indicators (equal to one if plan k was chosen) for all other employees in department j (not including employee i).[7]

The remainder of this section will assume for simplicity that the η_{jk} terms are zero (no common unobservables); I will return to discuss the issue of common unobservables in the next section. The idiosyncratic preference terms ϵ_{ik} are assumed to be extreme value deviates, so that the choice probabilities take the familiar logit form:

$$P_i(k|x_i;\beta) = \frac{exp\left\{x_i'\beta_k + f(d_{-i,jk})\right\}}{exp\left\{f(d_{-i,j1})\right\} + \sum_{m=2}^{K} exp\left\{x_i'\beta_m + f(d_{-i,jm})\right\}} \ , \tag{5}$$

where the coefficient vector β of the first plan is normalized to zero to resolve the usual indeterminacy. It remains to specify the function $f(\cdot)$. I will consider two specifications: one in which $f(d_{-i,jk})$ is just an unweighted share of coworkers choosing plan k, and another in which coworkers' influence is weighted according to their "proximity" to worker i in the space of characteristics.

4.1 All coworkers' choices have equal influence

Table 5 reports maximum likelihood estimates of the above model when $f(d_{-i,jk}) = \gamma \frac{1}{n_j} \sum_{l \neq i} d_{lk}$. In this case, employees' perceptions of the payoffs to a given plan increase in proportion to the share of their

[7] Note that the social effect looks like a traditional network effect: the utility of each choice depends on the decisions of others. (See Farrell and Saloner (1985) and Katz and Shapiro (1985) for seminal analyses of network effects.) *Direct* network externalities are probably unimportant in this context; however, indirect network effects may play a role if (for instance) popular plans are preferred because information about navigating their physician networks and coverage policies is more readily available.

coworkers who chose that plan, with γ indexing the strength or importance of the peer influence. For the results reported in the tables, the observable characteristics (x_i) include age, annual salary, tenure (years since hired), a sex dummy equal to one for males, coverage code dummies indicating whether the employee enrolled in a single, two-party, or family policy, and a full set of dummies for zipcodes of residence.[8] The estimates of γ should therefore be interpreted as the influence of coworkers' decisions after conditioning on the employee's own characteristics. Table 4 shows summary statistics by campus for the employee variables.

The estimates of the social effect γ are all positive and precisely estimated. Moreover, the estimates are strikingly similar across campuses, perhaps suggesting that the social learning process is similar in the separate locations. To get a sense for the magnitudes of the effects, consider two employees, both of whom are equally likely to choose Pacificare as Health Net based on their individual characteristics. If one of the employees is in a department where Pacificare's share is 10 percentage points higher than in the other employee's department, the former employee will be 19-22 percent more likely to choose Pacificare than the latter employee.

The table also reports the estimated coefficients on employee characteristics. Although these effects are not the focus of the present analysis, it is interesting to note that some of the effects are highly significant and consistent across campuses. For example, UC Care (the POS plan) is clearly preferred by employees that are older and have higher incomes. Kaiser is apparently much less attractive to females than males, and more attractive to single employees than employees with families.

One way to put the magnitude of the estimated social effect in perspective is to compare it with the influence of employees' demographic characteristics. For example, if the share of an employee's coworkers who choose UC Care goes up by 10 percentage points, this will increase the probability the employee chooses UC Care by roughly 20 percent. An equivalent impact on the choice probability would result from a 10-year increase in the employee's age, or a $5,000 increase in the employee's annual salary.

4.2 Closer neighbors have greater influence

The estimates of the social effect in the above specifications should be interpreted as the average influence of peers in one's department, where the average is taken over all employees in the department—including those with whom one never communicates. Especially in large departments, it is difficult to imagine that every employee in one's department is influential. Individuals are generally more likely to converse with peers that share common characteristics: young employees prefer to mingle with other young employees,

[8] The number of zipcode dummies ranges from 70-180, depending on the campus. Employees in "sparse" zipcodes (typically defined as having fewer than 15 total employees living in them) are grouped together into one "other" zipcode.

females often prefer to be social with other females, etc. If this is the case, an individual's choice of health plan should be influenced most by peers that have similar observable characteristics such as age, sex, and income. We can incorporate this notion into the discrete choice model by parameterizing a weighted social effect:

$$f(d_{-i,jk}) = \gamma \frac{1}{n_j} \sum_{l \neq i} w_{il} d_{lk} \ ,$$

where

$$w_{il} = exp\{-|z_i - z_l|' \rho\}$$

and z is a vector of demographic variables presumed to influence the degree of social interaction between two persons. Note that this function is maximal when all of employee i's coworkers chose plan k *and* all of those employees have employee i's same characteristics, since in that case $|z_i - z_l| = 0$ for all l and $w_{il} = exp\{0\} = 1$. The vector ρ determines the rate at which coworkers' decisions are downweighted as a function of demographic distance: for example, if one of the characteristics in z is age, then if the corresponding ρ is large, only the decisions of coworkers who are very close in age will have any influence. On the other hand, if ρ is zero, then all coworkers' choices have equal influence regardless of age differences.

Table 6 shows parameter estimates from this model for all five campuses. Not surprisingly, the estimate of γ increases dramatically, reflecting the fact that only close demographic neighbors exert a full social effect, and that the influence of such neighbors should be larger than the influence of the average coworker. The estimates suggest the influence of coworkers' decisions is downweighted when there are age or income gaps or gender differences. Not all of the weighting parameters (ρ) are precisely estimated for every campus; at UC Davis, for example, the weights on age and sex are statistically indistinguishable from zero. However, for every campus a likelihood ratio test rejects the hypothesis that the weighting parameters are jointly insignificant.[9] The point estimates are roughly similar across campuses, and the magnitudes are quite plausible. At UC San Diego, for instance, the decision of a coworker that is ten years older or younger receives 89 percent of the weight that a coworker of the same age receives; a $10,000 difference in annual salary would be weighted at 80 percent of the weight given to someone with the same salary; and decisions of coworkers of the opposite sex are weighted at 59 percent of the weight given to coworkers of the same sex. The decision of a coworker that is ten years older, $10,000 per year richer, and of the opposite sex is 42 percent as influential as the decision of a coworker who is identical in those three dimensions.

[9] The χ^2_3 test statistics (and the associated p-values) are reported in the table.

Note that estimating a model with social effects declining in demographic distance is similar in spirit to the tests proposed by other authors for separating the impact of social effects from the influence of common unobservable characteristics (i.e., the η_{jk} terms from equation 4). If these unobservables operate at the department level, they can be expected to mimic the effects estimated in the models where coworkers' decisions are unweighted. However, department-specific unobservables cannot mimic the patterns described here. In contrast, the data patterns appear to be fully consistent with observed patterns of social interaction (in particular, that individuals tend to socialize most with people that share common characteristics). In the next section I turn to some specific tests and robustness checks aimed at resolving the ambiguity of interpretation.

5 Social learning vs. common unobservables: some tests

The signs and magnitudes of the estimates reported above seem quite reasonable when interpreted as reflections of social learning or peer influence. However, there are important reasons to be cautious about settling on such an interpretation. The models of the previous section assumed away the "common unobservables" (η_{jk}) of equation 4. If in fact these effects are important, they will generate correlation in the decisions of employees within departments, and make it difficult to draw any clear inferences about the role of social learning. In the terminology of Manski (2000), the difficulty is distinguishing between "endogenous interactions" and "correlated effects."

As a starting point, it is useful to consider how important these correlated effects would have to be in order to generate the findings of the previous section. That is, suppose we assume the converse of what was assumed in section 4: endogenous interactions are unimportant ($f(d_{-i,jk}) = 0$), but employees within a department may share an idiosyncratic preference for certain plans ($\eta_{jk} \neq 0$). Moreover, suppose we think of the η_{jk}'s as normally distributed preference shocks that are independent across departments (and across plans within a department). Relative to the variance of the individual-specific preference shocks (ϵ), what would the variance of the department-specific preference shocks need to be in order to generate results like the ones reported above?

To answer this question, I generate simulated random utilities using

$$\hat{u}_{ijk} = x_i'\hat{\beta}_k + \sigma\eta_{jk} + \epsilon_{ik} \ ,$$

where $\hat{\beta}_k$ are the estimated plan-specific coefficients from the model in section 4.1, $\eta_{jk} \sim N(0,1)$, and ϵ_{ik} is a type I extreme-value deviate. These random utilities are then used to simulate plan choices, and the resulting "data" are used to re-estimate the model of section 4.1 to get an estimate of the "social effect" (γ). The variance parameter σ is chosen to calibrate the simulation so that the resulting $\hat{\gamma}$ is close to the one reported in table 5. These calibrations required σ's between 0.40 (for UCSD) and 0.49 (for UCLA) to match the previous estimates. Since the standard deviation of ϵ is 1.28, we can say that random department-specific preference shocks would need to have a standard deviation roughly one third as large as the standard deviation of individual-specific preference shocks in order to effectively mimic the results of the previous section. While not obviously unreasonable, the required standard deviations seem perhaps too large: in the context of health plan choices, one might expect idiosyncratic, employee-specific preferences to be far more important than any preferences held in common within departments.

The availability of data across different UC campuses permits a somewhat more direct test of the common unobservables hypothesis. If, for example, physicists at UCSD tend to cluster on a certain plan due to some unobservable characteristic η, one might expect physicists at other campuses to also share that characteristic (and also cluster on the same plan). In essence, the ¡same-department / different-campus¿ comparisons allow us to estimate the η's as fixed effects. Using only data on employees in academic departments that exist in three or more campuses (21 academic departments), and isolating attention to the four most popular plans, I estimated a simple multinomial logit model with employee characteristics and department fixed effects as explanatory variables. The results yield mixed evidence on the importance (or lack thereof) of the department-specific effects. A likelihood ratio soundly rejects the hypothesis that the department-specific effects are jointly zero ($\chi^2_{57} = 109.32, p = 0.000$), but only 7 of the 60 coefficients are individually significant.[10] Though in principle this cross-campus comparison is a promising way to distinguish between department-specific unobservables and social learning, in this case the test offers no clear-cut answers. The results don't suggest a salient role for common unobservables, but there are at least two reasons to expect that we wouldn't find them even if they were important. First, the sample sizes are small here, both because we must focus only on departments that exist at multiple campuses, and (more importantly) because there are only a handful of campuses, so that in essence each department-specific effect must be estimated using very few observations. Second, we may rightly be skeptical of the premise of this test: although it seems reasonable to expect, say, economists at UCLA to share similar unobservable characteristics with economists at Berkeley, the features of the choice set may differ in the two locations. For example, anecdotal evidence suggests that Kaiser is a much more attractive plan in Northern California than

[10]For example, the results indicate that chemists are significantly more likely to choose UC Care even after controlling for individual characteristics, and art historians are more likely to choose Pacificare. These and a few other department-specific effects are statistically significant; in none of these cases do I see an obvious explanation for the effect.

in Southern California. Therefore, employees of the same field might make different choices not because the "η's" are absent, but because the choice sets are only nominally identical.

An alternative way to get at the identification of social effects would be to employ tests similar to those used by Munshi and Myaux (2000) and Duflo and Saez (2000). In the present context, the idea is to break the sample into subgroups (within departments) that can plausibly be considered as the relevant social sub-networks. If the positive estimates of γ from the previous section reflect the influence of unobservable characteristics that are held in common among employees of the same department, then decisions of coworkers both inside *and* outside of an employee's subgroup (but still within the same department) will appear to be influential. If, on the other hand, department-specific unobservables are unimportant, only the own-group decisions should matter.

The obvious split to consider in my sample is that of faculty vs. staff within academic departments. To the extent that faculty socialize primarily with other faculty, and staffpersons socialize primarily with other staffpersons, we should find that only own-group social effects are important. Table 7 reports results from this exercise for the five UC campuses. As expected, the measured impact of peers' decisions within one's own subgroup is large and statistically significant; in contrast, the cross-group effects are smaller and (in all but two cases) indistinguishable from zero. For every campus except Davis, a likelihood-ratio test clearly rejects the hypothesis that own-group and cross-group effects are equal. These results are obviously consistent with the social learning hypothesis, and indicate that if common unobservables are important, they must not be held in common at the department level.[11]

Given the data available for the present study, a promising approach to distinguishing social learning from common unobservables is to look at the behavior of employees in special circumstances: in particular, those who have been newly hired and those who are switching plans. In the case of new hires, one might expect these employees to have relatively little private information about the available plans, and therefore to rely more heavily on the suggestions and observed choices of their coworkers. This hypothesis is easily tested within the empirical framework of the previous section. In particular, we can estimate the model of section 4.1 allowing the coefficient on department-level plan shares to differ for new vs. non-new employees:

[11] Although these findings favor the social learning interpretation of the data patterns, I'm inclined to be more cautious than previous authors in drawing this conclusion based solely on this test. The results suggest the effects of common unobservables are not important at the department level; however, such effects may still be operating at the subgroup-level. If the concern is that people who end up employed in economics departments share common unobserved characteristics that affect their (otherwise independent) health plan choices, then perhaps we should be doubly concerned that such effects also operate at the subgroup level. In other words, just because we have evidence that economics faculty don't share common unobservable characteristics with economics staff, we cannot necessarily infer that such characteristics aren't shared with other faculty.

$$u_{ijk} = x_i'\beta_k + [(1 - I_i^{new})\gamma_0 + I_i^{new}\gamma_1] \frac{1}{n_j} \sum_{l \neq i} d_{lk} + \epsilon_{ik} \ , \tag{6}$$

where I_i^{new} is an indicator equal to one if employee i was newly hired this period and ϵ_{ik} is still assumed to be double exponential. Table 8 presents maximum likelihood estimates of this model along with likelihood-ratio tests of the hypothesis that there is no difference between continuing and newly hired employees. The point estimates are all consistent with the learning explanation: the choices of newly hired employees appear to be more peer-influenced than the choices of continuing employees. The difference is statistically significant for Berkeley, Irvine, and San Diego, and insignificant for Davis and Los Angeles. Although the evidence isn't statistically clear for all five campuses, I find these results to be among the most compelling indications that our measured effects truly reflect social learning processes.

Finally, given the dynamics of individual learning, the decisions of employees who switch plans may yield additional evidence about the relevance of social learning. While new employees depend more heavily on peers' information and observed actions, the opposite may be true of employees who switch plans. In particular, plan switches may reflect decisions made based on new private information—e.g., if individuals "go with the crowd" until they obtain enough information to individually optimize. To the extent that this is a correct characterization of plan switches, we should observe that switches are "de-similarizing" in the sense that they are movements away from the clustered choices of the department as a whole. As a crude check of this hypothesis, table 9 lists the fraction of switches that reduced vs. increased the index of clustering (G) for the employee's department. Most of the switches were indeed "de-similarizing," though the majority is obviously narrow—and for two campuses the reverse was true. Given that health plan choices seemed to be more clustered at UCLA than at any other campus, it is interesting to note that at UCLA fully two thirds of the plan switches were desimilarizing switches.

6 Conclusions and Future Directions

Understanding the processes by which individuals choose health plans is of considerable interest to insurance executives, benefits coordinators, and healthcare policymakers. The various results of this study together suggest that individuals learn from and are influenced by their peers when making decisions about health insurance. More broadly, this is but one example of the many decisions made by economic agents in which learning from others plays a critical role in the decision-making process. The growing body of theoretical and empirical literature on social learning reflects a growing recognition of this role.

Ironically, though the importance of social learning is widely acknowledged in economics, social learning effects have been notoriously difficult to identify empirically. Some of the same identification problems that have plagued other studies are present here. In particular, the decisions of employees within departments appear to be highly correlated, but the difficult question is whether social learning can be distinguished from the influence of common unobservable characteristics as the source of this correlation. I offer a number of alternative tests aimed at resolving this issue. Taken together, the results of these tests broadly support interpreting the data patterns as evidence of social learning.

This study does not yet fully exploit the dynamic nature of the data: most of the analyses presented are cross-sectional in nature. Since learning (social or otherwise) is most realistically described as a dynamic process, it is natural to expect that some of the most interesting tests of social learning hypotheses will come from exploring dynamics in the data—for example, tracking employees' switches in health plans from year to year and examining the decisions of employees who switch departments or campuses within the sample period. Extensions of this nature will be the focus of future research.

References

Banerjee, A. V. (1992, August). A Simple Model of Herd Behavior. *Quarterly Journal of Economics 107*(3), 797–817.

Bernheim, B. D. (1994, October). A Theory of Conformity. *Journal of Political Economy 102*(5), 841–877.

Bertrand, M., E. Luttmer, and S. Mullainathan (2000, August). Network Effects and Welfare Cultures. *Quarterly Journal of Economics 115*(3), 1019–1055.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992, October). A Theory of Fads, Fashions, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy 100*, 992–1026.

Brock, W. A. and S. N. Durlauf (2001, April). Discrete Choice with Social Interactions. *Review of Economic Studies 68*(2), 235–260.

Buchmueller, T. C. and P. J. Feldstein (1997, April). The Effect of Price on Switching Among Health Plans. *Journal of Health Economics 16*(2), 231–47.

Conley, T. G. and C. R. Udry (2000). Learning About a New Technology: Pineapple in Ghana. Mimeo, Yale University.

Duflo, E. and E. Saez (2000, May). Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues' Choices.

Ellison, G. and D. Fudenberg (1993, August). Rules of Thumb for Social Learning. *Journal of Political Economy 101*(4), 612–643.

Ellison, G. and D. Fudenberg (1995, February). Word-of-Mouth Communication and Social Learning. *Quarterly Journal of Economics 110*(1), 93–125.

Ellison, G. and E. L. Glaeser (1997, October). Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach. *Journal of Political Economy 105*(5), 889–927.

Farrell, J. and G. Saloner (1985, Spring). Standardization, Compatability, and Innovation. *RAND Journal of Economics 16*(1), 70–83.

Glaeser, E. L., B. Sacerdote, and J. A. Scheinkman (1996, May). Crime and Social Interactions. *Quarterly Journal of Economics 111*(2), 507–548.

Hong, H., J. Kubik, and J. Stein (2001, July). Social Interaction and Stock Market Participation. *NBER Working Paper W8358*.

Katz, M. and C. Shapiro (1985, June). Network Externalities, Competition, and Compatibility. *American Economic Review 75*(3), 424–440.

Manski, C. F. (2000, Summer). Economic Analysis of Social Interactions. *Journal of Economic Perspectives 14*(3), 115–136.

Munshi, K. (2000, March). Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution. U. of Pennsylvania Mimeo.

Munshi, K. and J. Myaux (2000, July). Social Change and Individual Decisions: With an Application to the Demographic Transition. U. of Pennsylvania Mimeo.

Woittiez, I. and A. Kapteyn (1998, November). Social Interactions and Habit Formation in a Model of Female Labour Supply. *Journal of Public Economics 70*(2), 185–205.

Table 1: Plan Enrollments by Campus

	Number (percent) of employees enrolled				
Plan	**Berkeley**	**Davis**	**Irvine**	**Los Angeles**	**San Diego**
Health Net	4,469 (39.8)	5,524 (38.9)	1,778 (24.3)	5,272 (27.0)	4,571 (40.0)
Kaiser	3,774 (33.6)	3,112 (22.0)	1,566 (21.4)	4,021 (20.6)	2,335 (20.5)
Pacificare	1,313 (11.7)	2,386 (16.8)	1,542 (21.0)	2,459 (12.6)	2,146 (18.8)
UC Care	1,470 (13.1)	1,775 (12.5)	2,327 (31.7)	7,383 (37.9)	2,110 (18.5)
Prudential High	9 (0.1)	1 (0.0)	3 (0.0)	27 (0.1)	11 (0.1)
Core Medical	184 (1.6)	265 (1.9)	116 (1.6)	362 (1.9)	244 (2.1)
Western Health	—	1,142 (8.0)	—	—	—
Total	11,219 (100.0)	14,215 (100.0)	7,332 (100.0)	19,501 (100.0)	11,417 (100.0)

Table 2: Health plan clustering within departments

	Number (%) of departments with p-values less than:				
Campus	0.05	0.10	0.25	0.50	1.00
Berkeley	39 (18.3)	57 (26.8)	95 (44.6)	144(67.6)	213 (100)
Davis	44 (21.7)	60 (29.6)	102(50.2)	130(64.0)	203 (100)
Irvine	31 (23.1)	42 (31.3)	66 (49.3)	93 (69.4)	134 (100)
Los Angeles	101 (32.6)	133 (42.9)	198 (63.9)	243 (78.4)	310 (100)
San Diego	40 (22.1)	52 (28.7)	89 (49.2)	135(74.6)	181 (100)
Total	255 (24.5)	344 (33.0)	550 (52.8)	745 (71.6)	1,041 (100)

Table 3: Examples of predicted and actual plan shares for UCSD departments

	Economics		Linguistics		Physics		Cell. & Molec. Medicine	
Plan	Predicted	Actual	Predicted	Actual	Predicted	Actual	Predicted	Actual
Core	.01	.03	.01	.00	.01	.00	.03	.02
Pacificare	.15	.09	.18	.17	.15	.07	.23	.48
Health Net	.38	.41	.43	.46	.42	.39	.44	.35
Kaiser	.12	.13	.14	.13	.13	.11	.14	.12
UC Care	.33	.34	.23	.25	.30	.44	.16	.04
$E[G]$	.020		.027		.009		.013	
Observed G	.005		.002		.029		.088	
p-value	.863		.999		.023		.000	
# employees	37		24		80		53	

Table 4: Summary statistics by campus

	Berkeley	**Davis**	**Irvine**	**Los Angeles**	**San Diego**
Mean Age	44.4	43.5	42.3	41.9	42.6
(Std.Dev.)	(10.7)	(10.1)	(10.8)	(11.0)	(10.3)
Mean Income	51.6	48.4	50.3	52.5	49.0
(Std.Dev.)	(31.3)	(33.7)	(41.7)	(41.1)	(37.7)
Mean Tenure	10.2	9.3	8.4	8.1	8.6
(Std.Dev.)	(9.4)	(8.5)	(7.9)	(8.4)	(8.2)
Mean Sex	.461	.408	.379	.431	.418
(Std.Dev.)	(.499)	(.491)	(.485)	(.495)	(.493)
Employees	11745	14996	7686	20377	11886
Departments	407	389	298	442	284

Table 5: Maximum likelihood estimates, unweighted peer effects

	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Coworkers' share (γ)	1.721	1.830	1.698	1.999	1.673
	(0.109)	(0.120)	(0.130)	(0.065)	(0.106)
Health Net:					
Constant	0.140	-0.197	-0.574	-0.027	0.421
	(0.822)	(0.444)	(0.745)	(0.468)	(0.795)
Age	0.002	0.007	0.002	0.000	-0.000
	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
Salary	0.009	0.014	0.001	0.008	0.014
	(0.004)	(0.004)	(0.006)	(0.003)	(0.003)
Tenure	0.017	-0.000	0.061	0.065	0.042
	(0.005)	(0.005)	(0.009)	(0.006)	(0.005)
Male	0.023	0.089	0.032	0.124	0.066
	(0.071)	(0.070)	(0.102)	(0.057)	(0.061)
Two-party	-0.128	-0.193	0.121	-0.065	-0.059
	(0.090)	(0.093)	(0.128)	(0.075)	(0.077)
Family	-0.109	-0.319	-0.143	-0.272	-0.109
	(0.082)	(0.079)	(0.114)	(0.067)	(0.070)
Kaiser:					
Constant	-0.139	0.837	-1.113	-0.174	0.245
	(0.857)	(0.456)	(0.814)	(0.460)	(0.844)
Age	-0.005	-0.003	-0.001	-0.008	-0.006
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)
Salary	0.015	0.009	0.008	0.022	0.010
	(0.004)	(0.005)	(0.006)	(0.004)	(0.004)
Tenure	0.017	0.006	0.073	0.095	0.060
	(0.005)	(0.006)	(0.009)	(0.006)	(0.006)
Male	0.253	0.270	0.268	0.315	0.254
	(0.073)	(0.081)	(0.110)	(0.062)	(0.070)
Two-party	-0.071	-0.157	-0.027	-0.048	-0.005
	(0.092)	(0.104)	(0.139)	(0.082)	(0.090)
Family	-0.217	-0.535	-0.266	-0.193	-0.064
	(0.084)	(0.093)	(0.123)	(0.072)	(0.081)
UC Care:					
Constant	-1.718	-2.605	-4.215	-1.845	-1.249
	(0.944)	(0.616)	(1.209)	(0.471)	(0.867)
Age	0.018	0.023	0.018	0.016	0.014
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Salary	0.034	0.033	0.038	0.037	0.039
	(0.005)	(0.006)	(0.005)	(0.003)	(0.004)
Tenure	-0.012	-0.011	0.035	0.070	0.035
	(0.006)	(0.006)	(0.008)	(0.006)	(0.006)
Male	-0.030	0.100	-0.061	-0.162	0.013
	(0.091)	(0.093)	(0.101)	(0.059)	(0.076)
Two-party	-0.182	-0.390	0.152	-0.210	-0.194
	(0.112)	(0.116)	(0.123)	(0.076)	(0.094)
Family	-0.457	-0.819	-0.362	-0.614	-0.514
	(0.105)	(0.106)	(0.112)	(0.069)	(0.088)
N	9,896	8,004	4,239	15,690	10,019
Avg. Log-likelihood	-1.163	-1.344	-1.196	-1.119	-1.190

Standard errors in parentheses. Salary is annual salary in thousands; tenure is years since hired; two-party typically indicates coverage for employee plus spouse, and "family" indicates coverage for the employee plus two or more dependents. Coefficients for Western Health Advantage (UC Davis only) and for a full set of zipcode dummies are omitted to save space.

Table 6: Maximum likelihood estimates, weighted peer effects

	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Coworkers' choices (γ)	3.902	3.019	3.041	4.642	3.223
	(0.792)	(0.379)	(0.547)	(0.366)	(0.418)
Age weight (ρ_{age})	0.037	0.013	0.019	0.048	0.018
	(0.014)	(0.010)	(0.013)	(0.008)	(0.011)
Sex weight (ρ_{sex})	0.465	0.075	0.254	0.293	0.532
	(0.163)	(0.145)	(0.209)	(0.097)	(0.176)
Salary weight (ρ_{salary})	0.014	0.013	0.011	0.012	0.016
	(0.008)	(0.004)	(0.006)	(0.003)	(0.006)
N	9,896	8,004	4,239	15,690	10,019
Avg. Log-likelihood	-1.162	-1.342	-1.194	-1.115	-1.188
LR test of $\rho = 0$	20.66 (.000)	26.46 (.000)	25.84 (.000)	122.64 (.000)	32.80 (.000)

Standard errors in parentheses. Coefficients on other covariates are omitted to save space (the estimates are very similar to those reported in table 5). For the likelihood-ratio test of the hypothesis that the weight parameters (ρ) are all zero, the table reports the χ^2_3 test statistics with the corresponding p-values in parentheses.

Table 7: Academic departments divided into faculty and staff subgroups

	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Own group's share (γ_0)	1.599	1.084	1.189	1.383	1.304
	(0.163)	(0.239)	(0.221)	(0.169)	(0.263)
Other group's share (γ_1)	0.421	0.395	0.307	0.376	-0.264
	(0.157)	(0.233)	(0.227)	(0.164)	(0.270)
N	2,236	1,408	1,200	2,877	1,418
Avg. Log-likelihood	-1.157	-1.384	-1.162	-1.058	-1.089
LR test of $\gamma_0 = \gamma_1$	20.4 (.000)	1.2 (.279)	6.6 (.001)	16.7 (.000)	27.0 (.000)

Standard errors in parentheses. Plan-specific coefficients on employee characteristics are omitted to save space. For the likelihood-ratio test, χ^2_1 test statistics and p-values are shown.

Table 8: Differential social effects for new hires

	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Non-new employees (γ_0)	1.688	1.814	1.667	1.992	1.616
	(0.110)	(0.120)	(0.133)	(0.066)	(0.107)
New employees (γ_1)	2.671	2.325	2.359	2.193	3.679
	(0.418)	(0.465)	(0.586)	(0.288)	(0.455)
N	9,896	8,004	4,239	15,690	10,019
Avg. Log-likelihood	-1.163	-1.343	-1.188	-1.119	-1.189
LR test of $\gamma_0 = \gamma_1$	5.79 (.016)	1.23 (.267)	69.04 (.000)	0.48 (.488)	21.85 (.000)

Standard errors in parentheses. Plan-specific coefficients on employee characteristics are omitted to save space.
For the likelihood-ratio test, χ^2_1 test statistics and p-values are shown.

Table 9: Health plan switches and measures of clustering

	Number (%) of switches that:		
Campus	Decrease G	Increase G	Total
Berkeley	271 (57.1)	204 (42.9)	475
Davis	315 (42.8)	425 (57.4)	740
Los Angeles	557 (66.7)	278 (33.3)	835
San Diego	349 (58.1)	252 (41.9)	601
Irvine	193 (47.7)	212 (52.3)	405
Total	1,685 (55.1)	1,371 (44.9)	3,056

Figure 1: Distribution of G-percentiles: UC Berkeley

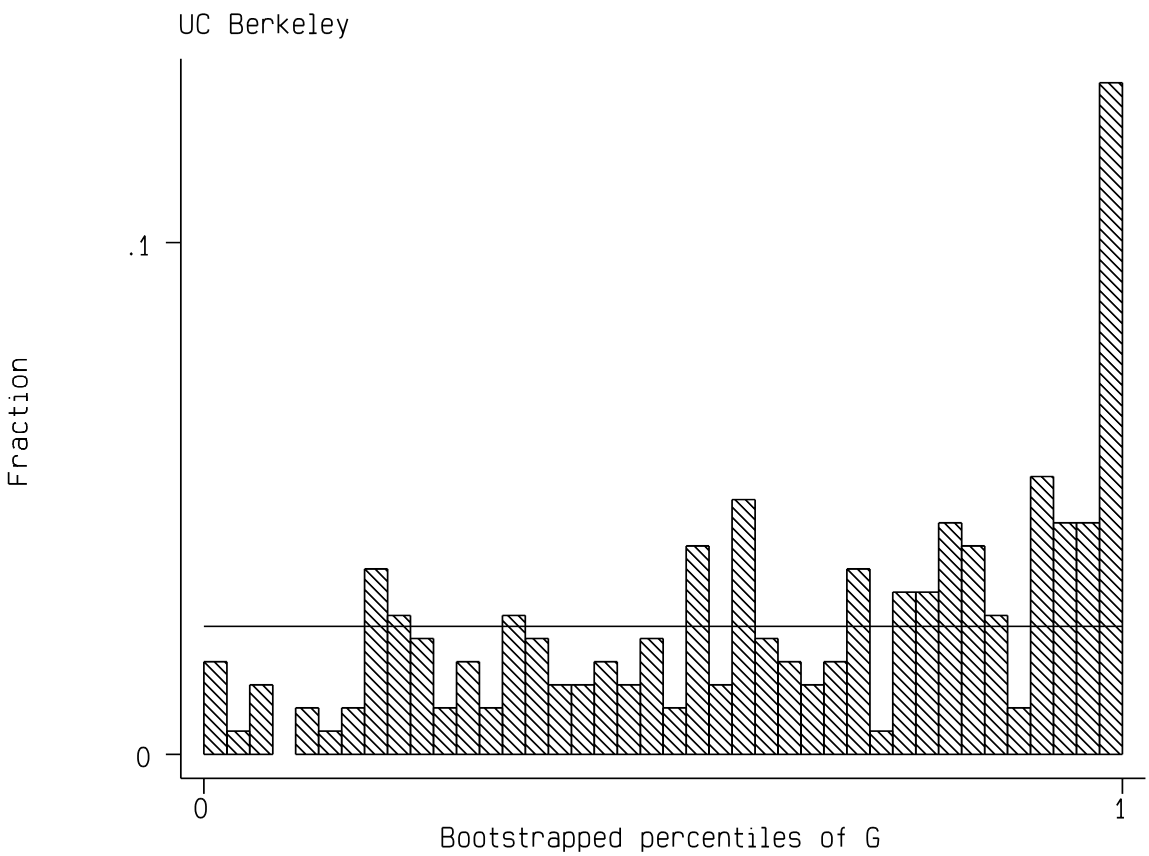

Figure 2: Distribution of G-percentiles: UC Davis

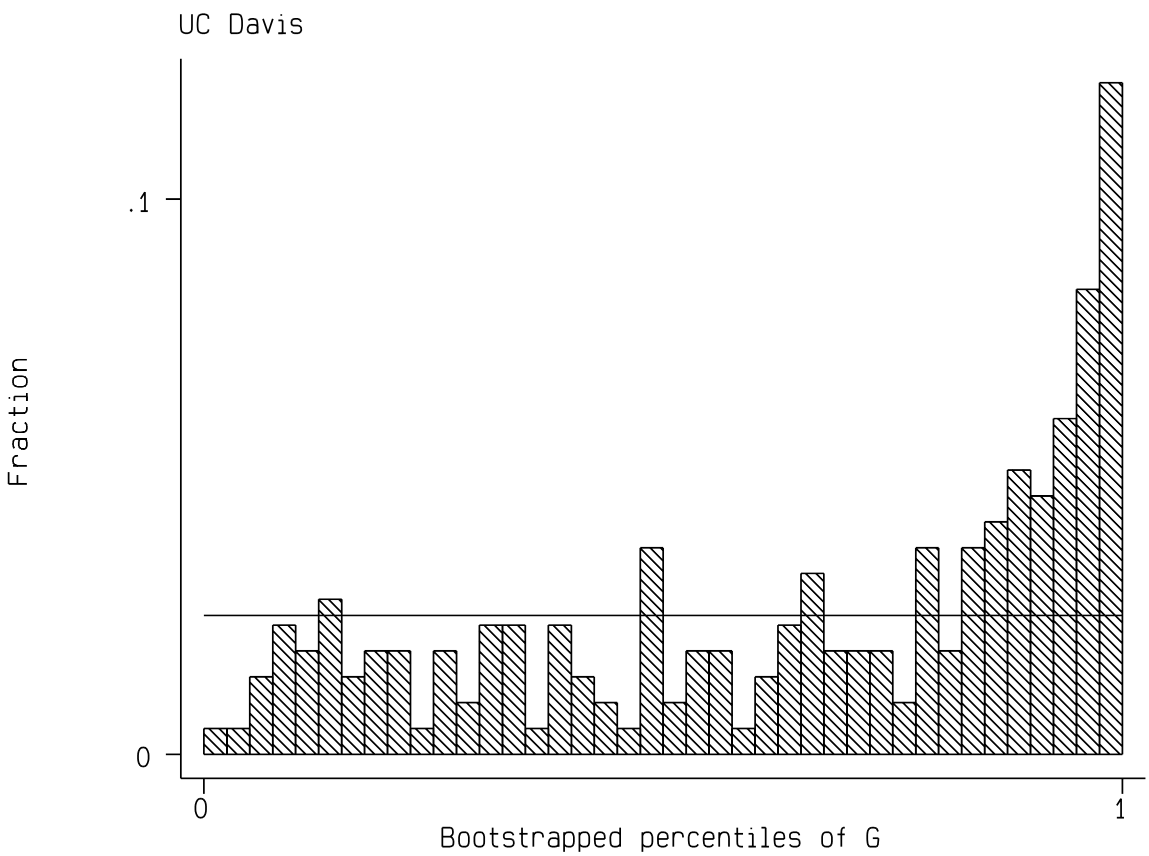

Figure 3: Distribution of G-percentiles: UC Irvine

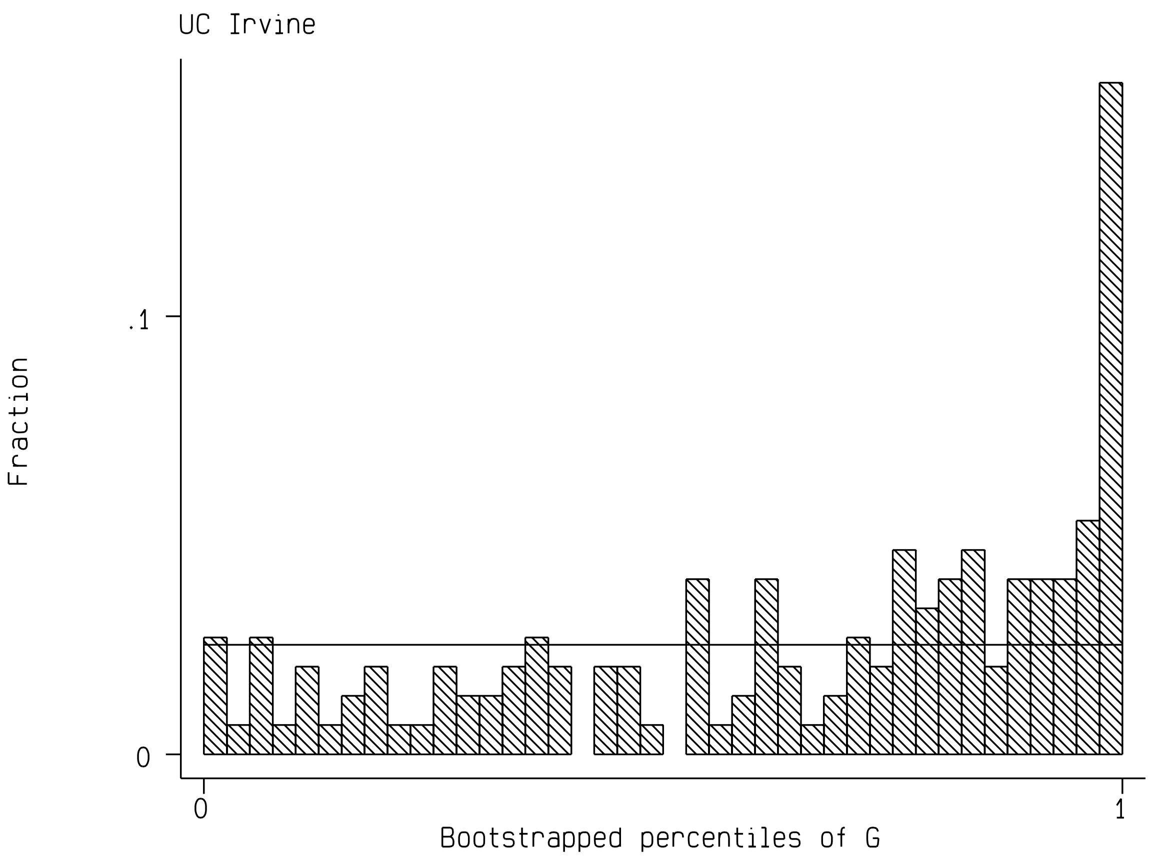

Figure 4: Distribution of G-percentiles: UC Los Angeles

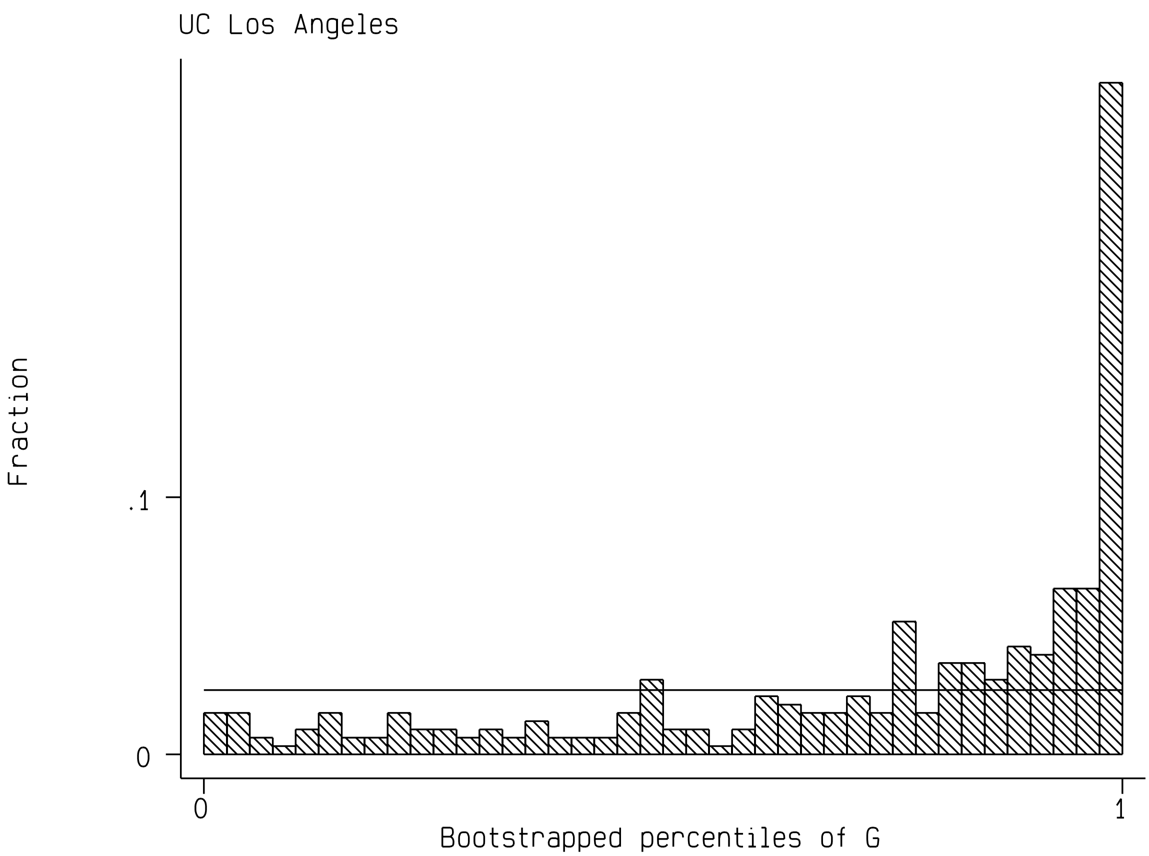

Figure 5: Distribution of G-percentiles: UC San Diego

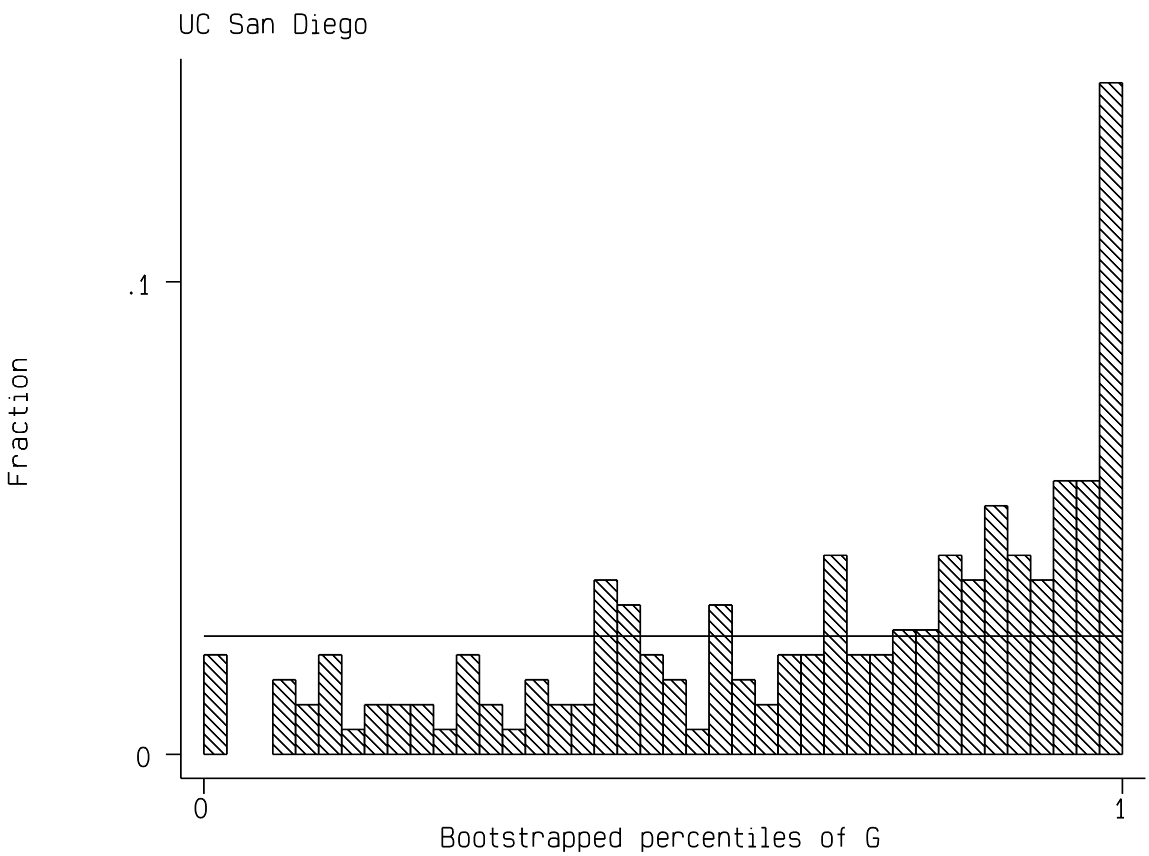

Figure 6: Distributions at UCSD with and without correction for demographics