

Social learning and health plan choice

Alan T. Sorensen*

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, health plan selections should appear to be correlated across employees within the same department. Estimates of discrete choice models in which individuals' perceived payoffs are influenced by coworkers' decisions reveal a significant (but not dominant) social effect. The strength of the effect depends on factors such as the department's size or the employee's demographic distance from her coworkers. The estimated effects are present even when the model allows for unobserved, department-specific heterogeneity in employee preferences, so the results cannot be explained away by unobservable characteristics that are common to employees of the same department.

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 parallel decisions involving uncertain payoffs.

The objective of this article is to quantify the impact of social learning in a specific context: individuals' choices of employer-sponsored health plans. Using data from five University of California campuses, I show that health plan choices are “clustered” within departments: the choices of employees in the same department are too similar relative to what we would expect based on individual characteristics and campus-wide patterns. I estimate econometric models of health plan choice that explicitly allow for social learning and find large, statistically significant effects that are robust across campuses and model specifications.

Health care market performance is a matter of considerable policy importance, and understanding demand in health care markets is essential to credible policy analysis. In the demand for employer-sponsored health insurance, uncertainty about product quality plays a critical role. As with other experience goods, the quality of a health plan (or one's match quality with a health plan) is very difficult to know *ex ante*, and choosing the “wrong” plan can be costly—especially

* Stanford University and NBER; asorensen@stanford.edu.

I thank Michele French and the University of California for providing the data, Zhigang Li for outstanding research assistance, the Editor and referees for helpful comments, and the National Science Foundation for financial support (grant no. SES-0079201). Any errors are mine.

because the costs of switching can be high. This suggests that providing information—which employers generally do by distributing brochures or hosting benefits seminars—is critical. But in some cases the most valuable information comes from others who have already experienced the experience good. By assessing the significance of peer influence in employees' health plan selections, this article sheds some light on the process through which information diffuses in a large organization, and it also speaks to the importance of peer recommendations in the demand for experience goods more generally.

While the findings on health plan enrollment decisions are interesting in their own right, this article's contribution to the broader literature on social learning extends beyond the specific context of health plan choice. A growing body of theoretical research has incorporated social learning into standard models of economic decision making, showing that social effects can alter those models' predictions in important ways.¹ However, social effects have been notoriously difficult to quantify empirically, in large part due to identification problems that have been described in detail by Manski (2000). Of principal relevance 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. Previous empirical studies have attempted to resolve this issue in a variety of ways, with varying degrees of success.

Relative to the existing empirical literature, the contribution of this article is to analyze a dataset in which the presence of social learning can be demonstrated convincingly even while directly allowing for unobserved heterogeneity. The data permit a rich set of comparisons and tests that distinguish social learning from correlation in unobserved preferences. Most directly, the panel nature of the data allows me to simultaneously estimate social effects and department-specific unobservables. The estimated social effects are somewhat smaller but remain significant when department-specific unobservables are included. A different robustness check (similar in spirit to Munshi and Myaux (2000)) examines the own- and cross-group influences among faculty versus staff, finding that own-group effects are large and statistically significant, while cross-group effects are in most cases indistinguishable from zero. Moreover, the own-group effects are no larger among faculty than among staff, even though the common unobservables problem is probably more relevant to faculty.

Another well-known difficulty of estimating models with social interactions is that observed choices are jointly endogenous: coworkers' choices cannot be regarded as exogenous influences, since they are in turn influenced by the choice of the employee in question. Rather than incorporating some notion of equilibrium in the estimation, as some authors have suggested, I handle this issue by focusing solely on newly hired employees: I assume that new hires are potentially influenced by the observed choices of existing employees, but not vice versa. (I explain in Section 3 why this is a reasonable assumption.) This approach has the obvious advantage of simplifying the estimation problem, and it also focuses the model's attention on employees who are actively making health plan decisions.

Social-learning hypotheses have been studied previously in many other contexts, including crime (Glaeser, Sacerdote, and Scheinkman, 1996), labor supply (Woittiez and Kapteyn, 1998), contraception (Munshi and Myaux, 2000), adoption of fertilization technology (Conley and Udry, 2000; Munshi, 2000), welfare program participation (Bertrand, Luttmer, and Mullainathan, 2000), stock market participation (Hong, Kubik, and Stein, 2001), and labor market outcomes (Bayer, Ross, and Topa, 2005). The study most similar to this one is that of Duflo and Saez (2002), who examine individuals' decisions about whether to enroll in a university-sponsored retirement plan. Although similar in spirit to the present analysis, the instrumental variables approach used by Duflo and Saez is fundamentally different from the identification strategies employed here, which

¹ See, for example, the articles by Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Ellison and Fudenberg (1993).

rely primarily on data variation over time and across campuses.² Also, this article specifically addresses the issue of quantifying social effects in a polychotomous discrete choice setting.

2. Background and data

■ The University of California (UC) system is made up of 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),³ 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 extremely 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.⁴ Plan enrollments by campus are shown in Table 1 for the year 2000.

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 nonprice considerations, since the employee premiums are essentially the same (zero in most years), and the required copays (e.g., for office visits or prescription drugs) are also generally uniform across plans. 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.⁵

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 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 staff person assigned as the benefits coordinator who may serve as the point person for information within departments. The information provided is intentionally neutral (in the sense that it doesn’t favor any particular plan) because presumably

² Their study utilizes multiple observations on the same individual over time, and uses an instrumental variables strategy suggested by Case and Katz (1991) to address the endogeneity problem potentially induced by common unobservables. The salary or tenure composition of the individual’s department is used as an instrument for changes in coworkers’ participation rates. In contrast, my article uses only the (one-time) decisions of new employees, and it regards movements over time in a department’s average choices as a direct source of identifying variation.

³ One additional HMO, Western Health, was available at UC Davis.

⁴ This default option only covers catastrophic medical care, so virtually all employees choose another plan (as shown in Table 1). The few who remain in the Core Medical plan are typically employees who are ineligible for one of the other more comprehensive plans (e.g., because they work less than half time).

⁵ A substantial change in the pricing structure occurred just prior to my 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.

TABLE 1 Plan Enrollments by Campus (2000)

Plan	Number (Percentage) of Employees Enrolled				
	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 (.1)	1 (.0)	3 (.0)	27 (.1)	11 (.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)

different plans will be optimal for different people. The plan shares shown in Table 1 seem to confirm the idea that no single option is uniformly superior.⁶

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 doctors, the incidence of reimbursement hassles, the difficulty of getting access to specialists, etc.—is not disseminated systematically. As a consequence, much of the information used as an input to the health plan decision is 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. This is the fundamental idea of the empirical analyses in the following sections: essentially, if we control for similarity in the observed characteristics of employees in the same department, and also for time-invariant unobserved characteristics shared by those employees, then the degree of excess similarity in their choices reveals the strength of the social-learning effect.

□ **The role of social learning.** In the existing literature, the term “peer effects” generally refers to a broad class of situations in which the action of an individual varies with the actions or characteristics of his peers. Such effects can arise for a variety of reasons, not all of which involve social interactions. The “social learning” examined in this article should be considered a specific form of endogenous peer effect, in which individuals are directly influenced by knowledge of their peers' decisions, or by information provided by peers. I ignore the possibility of exogenous peer effects (what Manski calls contextual effects) in the analysis, because it seems highly unlikely that individuals' health plan selections are influenced by the exogenous characteristics (such as age and sex) of their coworkers.

In the context of health plan choices, social learning can take a variety of different forms, all of which are likely to cause peers' decisions to be positively correlated. The simplest example occurs when individuals, facing uncertainty about the relative payoffs of the available options, draw inferences about plan qualities by observing the choices of their coworkers.⁷ 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.,

⁶ This feature of health plan decisions is an interesting contrast to decisions about retirement benefits, where, for example, it can be argued that every employee should contribute to a 401(k) plan.

⁷ This observational learning is an example of what Manski (2000) would call an “expectations interaction.”

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.⁸

One simple measure of the degree of similarity among coworkers' health plan decisions is

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

where k indexes health plans, s_{jk} is the share of department j 's employees who chose plan k , and x_k is the share of all campus employees who chose plan k .⁹ G_j is maximal if everyone in department j chooses the same plan, and minimal (zero) when the pattern of health plan choices in department j exactly matches the average pattern for the whole university. Calculating G for the departments in the sample reveals considerable clustering: nearly half (47.8%) of all departments have values of G that fall in the top decile of the appropriate sampling distribution, and nearly a quarter (23.7%) are in the top percentile.¹⁰ In other words, the observed values of G are generally too large to have resulted from choices that were made independently.

However, note that measuring x_k in equation (1) as the share of plan k in the entire university ignores differences in demographic composition across departments. 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 ideally be compared to appropriate, department-specific benchmarks that account for the composition of employee characteristics. In a previous version of this article I outline a method for making this kind of controlled comparison and show that even after controlling for departments' demographic compositions, there is still considerable excess clustering.

3. Empirical models and results

■ Under the simple assumption that individuals communicate with coworkers within their departments, the observed correlation in health plan choices across employees within a department should be seen as consistent with social learning. Of course, 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 in a later subsection.

Moving beyond basic correlations, the empirical objective of this article is to measure the impact of coworkers' choices on an individual's choice, and to compare that measured impact to the impact of other relevant factors. As a starting point, consider a model in which employee i in department j gets random utility from choosing plan k equal to

$$u_{ijk} = x_i' \beta_k + f(d_{-i,jk}) + \eta_{jk} + \varepsilon_{ik}, \quad (2)$$

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

⁸ 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. Moreover, 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.

⁹ Ellison and Glaeser (1997) used this measure as a starting point for analyzing industry agglomeration.

¹⁰ I simulated the sampling distribution of G under the assumption that individuals' choices were independent draws from a multinomial distribution, with probabilities given by the campus-wide plan shares (x_k). Since the distribution depends on the number of employees, percentiles were obtained separately for each department.

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).¹¹

Inspecting equation (2), one is immediately confronted with the two main difficulties of estimating an empirical model with social effects: simultaneity and unobserved heterogeneity. Regarding simultaneity, notice that individual i 's perceived payoffs are influenced by individual j 's choices, while individual j 's choice is likewise influenced by i 's decision. Intuitively, to the extent this simultaneity is important, it may tend to generate upward-biased estimates of the social effect, since the $d_{-i,k}$ terms (coworkers' choices) will be positively correlated with ε_{ik} . More fundamentally, however, the interdependence of coworkers' decisions means that estimation by ordinary maximum-likelihood methods is in principle inappropriate: the likelihood function cannot be specified as the product of individual likelihoods, and the joint likelihood (for the vector of employees' choices in a given department) cannot be obtained because the equilibrium of the simultaneous choice game may not be unique. The choice process implied by equation (2) can clearly have multiple equilibria, especially if the social effects are large.¹²

While various econometric strategies have been proposed for handling the simultaneity inherent in social effects models,¹³ the present data offer a natural and novel way to circumvent the potential problems: in estimating the choice models implied by equation (2), I can focus solely on the choices of newly hired employees. In other words, I can assume that health plan choices of newly hired employees are influenced by coworkers, but not vice versa, so that coworkers' decisions (the $d_{-i,k}$ terms) can be taken as exogenous with respect to a new employee's choice. This is a strong but perfectly reasonable assumption in this particular context: unlike existing employees, new employees' plan choices must reflect the outcome of an active decision process (as opposed to a passive, stay-in-last-year's-plan decision). Plan switches are very rare (the aggregate switch rate is approximately 5–6% in any given year), which means that the observed choices of existing employees largely reflect decisions that were made in the past. Moreover, even for existing employees who choose to switch plans, such decisions are unlikely to be influenced by the choices or suggestions of new employees, since new employees will in general have little information or experience concerning the plans.¹⁴

Of course, focusing the estimation on newly hired employees does not resolve the second difficulty inherent in the model: unobserved heterogeneity. In particular, the presence of department-specific preferences that are unobservable to the econometrician—the η_{jk} 's—makes it difficult to identify social effects other than through functional-form assumptions.¹⁵ Intuitively, if we observe that employees in a department are heavily clustered in Pacificare, the clustering can be attributed either to social effects or to unobserved, department-specific preferences; without any additional identifying variation in the data, one cannot distinguish the two explanations.

In the present study, various features of the data make it possible to test directly for the importance of department-specific unobservables (the η_{jk} 's). Most important, the panel nature of the data makes it possible to estimate the η_{jk} 's directly for a subset of the departments in the sample.

¹¹ 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 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 readily available from peers.

¹² For instance, if learning from peers completely dominates all other considerations, "everyone chooses Pacificare" and "everyone chooses Health Net" may both be equilibria.

¹³ Brock and Durlauf (2001) describe estimation methods based on equilibrium choices. Tamer (2000) offers a useful discussion of the econometric difficulties associated with multiple equilibria, along with an estimation approach that accounts for multiple equilibria.

¹⁴ Switching patterns in the data are consistent with this assumption: average switch rates for departments with no new hires in a given year are essentially the same as for departments with one or more new hires (5.58% versus 5.50%, respectively). Also, in probit models predicting plan switches, the number of newly hired employees in one's department has no statistically discernible impact.

¹⁵ This problem has been thoroughly explicated by Manski (2000), for example.

For simplicity and ease of exposition, the remainder of this section will assume that unobserved heterogeneity is unimportant—i.e., the η_{jk} 's are zero although in the final subsection I shall return to discuss the issue in depth. Also for the sake of simplicity, I assume the idiosyncratic preference terms ε_{ik} are extreme value deviates, so that the choice probabilities take the familiar logit form:

$$P_i(k | x_i; \beta) = \frac{\exp\{x_i' \beta_k + f(d_{-i,jk})\}}{\exp\{f(d_{-i,j1})\} + \sum_{m=2}^K \exp\{x_i' \beta_m + f(d_{-i,jm})\}}, \quad (3)$$

where the coefficient vector β of the first plan is normalized to zero to resolve the usual indeterminacy.¹⁶ It remains to specify the function $f(\cdot)$. I will consider three specifications: one in which $f(d_{-i,jk})$ is just an unweighted share of coworkers choosing plan k , another in which coworkers' influence is weighted according to their "proximity" to worker i in the space of characteristics, and another in which the per-coworker influence depends on the size of the employee's department.

□ **Baseline model: all coworkers' choices have equal influence.** Table 3 reports maximum-likelihood estimates of the model when $f(d_{-i,jk}) = \gamma(1/n_j) \sum_{\ell \neq i} d_{\ell k}$. In this case, employees' perceptions of the payoffs to a given plan increase in proportion to the share of their 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, a sex dummy equal to one for females, coverage code dummies indicating whether the employee enrolled in a single, two-party, or family policy, and a full set of dummies for region or neighborhood of residence; Table 2 shows summary statistics by campus for these variables.¹⁷ Including these variables as controls means that the estimates of γ should be interpreted as the influence of coworkers' decisions after conditioning on the employee's own characteristics.

The estimates of the social effect γ are all positive and precisely estimated. Moreover, the estimates are 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 as 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 14–21% more likely to choose Pacificare than the latter employee (depending on the campus).

The table also reports the estimated coefficients on employee characteristics. Although these effects are not the focus of the present analysis, it is useful 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 who are older and have higher incomes. Kaiser is apparently much less attractive to males than females, and typically 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. At UCLA, 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%. An equivalent impact on the choice probability would result from a six-year increase in the employee's age or a \$9,000 increase in the employee's annual salary.

□ **Weighted social effects: closer neighbors have greater influence.** The estimates of the social effect reported in Table 3 reflect the average influence of peers in one's department, where

¹⁶ In a recent study of crop-sharing contracts in Illinois, Young and Burke (2001) analyze a dynamic model in which the choice probabilities take this form. The social influence function, $f(\cdot)$, corresponds to their "local conformity" effect.

¹⁷ The number of region dummies ranges from 5 to 9, depending on the campus. Each region is defined as a group of adjacent zip codes, with the groupings constructed as naturally as possible (e.g., based on geographic boundaries). Aggregating zip codes greatly reduces the number of parameters to be estimated and has little impact on the estimates of the parameters of interest.

TABLE 2 Summary Statistics by Campus (2000)

	Berkeley	Davis	Irvine	Los Angeles	San Diego
Mean age	44.4 (10.7)	43.5 (10.1)	42.3 (10.8)	41.9 (11.0)	42.6 (10.3)
Mean income	51.6 (31.3)	48.4 (33.7)	50.3 (41.7)	52.5 (41.1)	49.0 (37.7)
Mean tenure	10.2 (9.4)	9.3 (8.5)	8.4 (7.9)	8.1 (8.4)	8.6 (8.2)
Mean female	.461 (.499)	.408 (.491)	.379 (.485)	.431 (.495)	.418 (.493)
Employees	11,745	14,996	7,686	20,377	11,886
Departments	407	389	298	442	284

Standard deviations in parentheses.

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 who share common characteristics: young employees prefer to mingle with other young employees, females may 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 who 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:

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

where

$$w_{i\ell} = \exp\{-|z_i - z_\ell|' \rho\}$$

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 $|z_i - z_\ell| = 0$ for all ℓ and $w_{i\ell} = \exp\{0\} = 1$. 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 4 shows parameter estimates from this model for all five campuses. Not surprisingly, in every case the estimate of γ increases relative to Table 3, 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 point estimates suggest that the influence of coworkers' decisions is downweighted when there are age or income gaps, or when there are differences in sex or city of residence. Many of the individual weighting parameters (ρ) are statistically indistinguishable from zero, but for every campus a likelihood ratio test rejects the hypothesis that all the weighting parameters are jointly zero.

The specification reported in the table allows the impact of age and salary differences to be asymmetric—i.e., younger or lower-ranked coworkers may be downweighted more than older, higher-ranked coworkers. The point estimates suggest hierarchical effects with respect to salary: at every campus, coworkers with higher salaries are downweighted less than coworkers with lower salaries. The coefficients on age differences are somewhat less precise, but at campuses where the difference is statistically significant the implication is that younger coworkers are more influential

TABLE 3 Maximum-Likelihood Estimates, Unweighted Peer Effects

Plan	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Coworkers' share (γ)	1.605 (.181)	1.872 (.183)	1.336 (.149)	1.859 (.088)	2.031 (.168)
Health Net:					
Constant	.087 (.197)	-.024 (.177)	-.588 (.265)	-.137 (.147)	-.639 (.154)
Age	.014 (.005)	.005 (.004)	.003 (.005)	.008 (.004)	.010 (.004)
Salary	-.003 (.003)	-.001 (.002)	.001 (.003)	-.003 (.002)	.001 (.002)
Female	.020 (.091)	-.054 (.091)	-.034 (.121)	.125 (.063)	.114 (.074)
Two-party	-.258 (.122)	-.161 (.092)	-.274 (.130)	-.074 (.075)	-.085 (.092)
Family	-.195 (.121)	-.100 (.081)	-.291 (.131)	-.205 (.070)	.049 (.082)
Kaiser:					
Constant	.403 (.195)	.190 (.185)	.179 (.274)	.263 (.161)	-.353 (.180)
Age	.016 (.005)	.005 (.004)	.000 (.006)	.006 (.004)	.006 (.005)
Salary	-.012 (.003)	-.013 (.003)	-.006 (.004)	-.017 (.003)	-.011 (.003)
Female	.216 (.096)	.250 (.076)	.164 (.111)	.236 (.066)	.148 (.086)
Two-party	-.356 (.121)	-.287 (.108)	-.134 (.164)	.105 (.083)	-.088 (.110)
Family	-.214 (.127)	-.339 (.087)	-.295 (.109)	.008 (.073)	.032 (.099)
UC Care:					
Constant	-2.551 (.233)	-2.504 (.188)	-3.097 (.317)	-1.816 (.146)	-2.647 (.187)
Age	.041 (.006)	.036 (.005)	.049 (.005)	.033 (.004)	.037 (.005)
Salary	.018 (.003)	.018 (.002)	.015 (.002)	.020 (.002)	.018 (.002)
Female	-.013 (.113)	.116 (.093)	-.233 (.113)	-.109 (.066)	.263 (.090)
Two-party	-.129 (.138)	-.250 (.112)	-.225 (.121)	-.237 (.077)	-.307 (.121)
Family	-.541 (.154)	-.798 (.126)	-.756 (.106)	-.629 (.074)	-.582 (.104)
N	6,345	6,911	3,780	13,403	6,319
Average log-likelihood	-1.165	-1.289	-1.252	-1.204	-1.210

Note: Estimates of baseline model, using only choices of newly hired employees. Standard errors in parentheses, adjusted for clustering by department. "Salary" is annual salary in thousands; "two-party" typically indicates coverage for employee plus spouse; "family" indicates coverage for the employee plus two or more dependents. Coefficients for Western Health Advantage (UC Davis only) and for full sets of regional dummies omitted to save space.

than older coworkers. This pattern could reflect that on average, younger employees made their health plan choices more recently, so the information they can share is more current. In general, the point estimates of the weighting coefficients are similar across campuses, and the magnitudes are plausible. At UC Berkeley, for instance, the decision of a coworker who is ten years older

TABLE 4 Maximum-Likelihood Estimates, Weighted Peer Effects

Plan	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Coworkers' choices (γ)	4.995 (1.192)	3.536 (.515)	4.627 (1.285)	4.583 (.442)	4.121 (.690)
Distance weights (ρ):					
Age (higher)	.036 (.035)	.006 (.018)	.025 (.030)	.041 (.009)	.040 (.022)
Age (lower)	.022 (.013)	.016 (.010)	.028 (.017)	.031 (.006)	.020 (.010)
Salary (higher)	.023 (.009)	.000 (.003)	.030 (.013)	.007 (.002)	.012 (.005)
Salary (lower)	.081 (.035)	.019 (.011)	.060 (.051)	.027 (.008)	.042 (.024)
Opposite sex	.320 (.183)	.123 (.163)	-.268 (.253)	.096 (.090)	.133 (.153)
Different city	.423 (.164)	.643 (.157)	.702 (.243)	.418 (.098)	.265 (.179)
<i>N</i>	6,345	6,911	3,780	13,403	6,319
Average log-likelihood	-1.1443	-1.2874	-1.2482	-1.2000	-1.2084
LR test of $\rho = 0$	262.7 (.000)	22.12 (.001)	28.73 (.000)	107.2 (.000)	2.22 (.003)

Note: Standard errors in parentheses. Coefficients on other covariates omitted to save space (estimates are very similar to those reported in Table 3). For likelihood-ratio test of hypothesis that weight parameters (ρ) are all zero, the table reports χ^2 test statistics with corresponding *p*-values in parentheses.

receives 70% of the weight given to a coworker of the same age; the decision of a coworker whose salary is \$10,000 higher would receive 79% of the weight given to a coworker with the same salary; and the decision of a coworker of the opposite sex receives 73% of the weight given to a coworker of the same sex. The decision of a coworker who is ten years older, \$10,000 per year richer, and of the opposite sex is 40% as influential as the decision of a coworker who is identical in those three dimensions.

The estimates suggest that individuals are influenced more by coworkers who live in the same city than by coworkers living in different cities.¹⁸ While this is consistent with the notion that decisions of coworkers in one's neighborhood are more informative (or that those coworkers have more relevant information to share), it is also consistent with the idea that individuals first choose a doctor in their home neighborhood and then choose a plan that covers that doctor—i.e., correlation in the decisions among individuals living in the same town could arise from unobserved heterogeneity rather than social interactions.

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 unobservables (i.e., the η_{jk} terms from equation (2)), such as the one proposed by Munshi and Myaux (2000) in which own-group effects are compared with cross-group effects.¹⁹ The idea is that estimated effects should get stronger when the social network is defined more narrowly. However, the example of different-city effects shown in Table 4 highlights the problem with such tests: just as we expect social interactions to be more salient in a more narrowly defined

¹⁸ In the reported specification, cities are defined somewhat broadly as collections of roughly 10–15 neighboring zip codes. So, for example, Emeryville and northern parts of Oakland are included with Berkeley, but Alameda and San Leandro are not.

¹⁹ Munshi and Myaux 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.

social group, we can also expect unobservable characteristics to be more highly correlated, so looking within subgroups or weighting peer effects by demographic distance cannot fully resolve the question of common unobservables. (Below I turn to some specific tests and robustness checks aimed at resolving the ambiguity of interpretation.) Nevertheless, these kinds of models can reveal the level at which common unobservables must operate in order to be a viable alternative explanation. For example, while department-specific unobservables could generate the unweighted results in Table 3, they could not be expected to mimic the results in Table 4: the unobservable characteristics would need to be specific to narrower subgroups (within departments) based on age, salary, sex, and home city.

□ **The influence of department size.** In the above specifications, the strength of the social effect is *de facto* scaled to be independent of department size: the impact of coworkers' choices depends only on the share of coworkers in each plan, not on the total number of coworkers. Given that departments in the sample vary considerably in size, an interesting question is whether this restriction is valid. If 6 out of 10 coworkers chose Health Net, does that have the same impact on perceived utility as 60 out of 100 coworkers choosing Health Net?

Especially in large departments, employees do not generally communicate with every one of their coworkers. A reasonable assumption is that individuals make regular contact with a small group of coworkers—say, 10 to 20—regardless of the overall size of the department. Therefore, an individual working in a 15-person department may communicate regularly with every coworker, whereas an individual working in a 100-person department communicates with only a fraction of his coworkers. The average per-coworker influence will therefore be lower for the employee in the larger department.²⁰

To examine this hypothesis, consider estimating the discrete choice model with the social effects entering as follows:

$$f(d_{-i,k}) = \sum_{i \neq j} \gamma_{ij}(N) d_{jk},$$

with $\gamma(N) = \gamma_0 + \gamma_1 N + \gamma_2/N$ representing the average per-coworker influence.²¹ Table 5 reports results from estimating such a model. As expected, the per-coworker influence declines with department size: the expected influence of a randomly selected coworker is larger in small departments. What may be unexpected is the precision of the estimates (most of the γ coefficients at all the campuses are significant at the 1% level) and the remarkable similarity of the estimates across campuses. Figure 1 plots the estimated per-coworker influence as a function of department size for each of the five campuses separately, and it is difficult to distinguish the five functions. An interesting avenue for future research will be to identify the kinds of social interaction that could give rise to such a consistent pattern.

□ **How big are the social effects?** To put the estimated social effects into perspective, it is useful to compare their importance in individuals' utility functions to the importance of other individual characteristics. If the variance across plans in social factors is much smaller (larger) than the across-plans variance in individual-specific preference factors, this would indicate that social effects play a relatively unimportant (important) role in determining individuals' choices. Stated in terms of the utility function in equation (2), if $x'_i \beta_k$ ranges from -8 to 8 across plans but $f(d_{-i,jk})$ ranges only from 0 to 2 , then individual demographics (the x 's) would clearly be the more important drivers of individuals' decisions.

²⁰ There are other reasons why per-coworker influences may appear to be smaller in large departments. If large social networks naturally gravitate toward being collections of smaller "cliques," then clustering at the department level may be difficult to discern even if clustering is extreme within cliques. (Imagine a large department as a collection of several cliques, with members in the same clique all making identical choices, but choices being independent across cliques. If the number of independent cliques is large enough, choices will appear unclustered at the department level.)

²¹ Note that this can equivalently be interpreted as $\gamma(N)$ times the share of coworkers in plan k , with $\gamma(N)$ being a simple quadratic function of N .

TABLE 5 Social Effects and Department Size

Plan	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
$\hat{\gamma}_0$.0136 (.0028)	.0025 (.0007)	.0013 (.0033)	.0004 (.0007)	.0081 (.0021)
$\hat{\gamma}_1$	-3.2×10^{-5} (7.0×10^{-6})	-6.9×10^{-7} (2.1×10^{-7})	-3.7×10^{-7} (1.9×10^{-6})	-5.3×10^{-7} (6.8×10^{-7})	-1.0×10^{-5} (4.4×10^{-6})
$\hat{\gamma}_2$	1.133 (.152)	1.668 (.140)	1.221 (.171)	1.814 (.078)	1.438 (.158)
<i>N</i>	6,345	6,911	3,780	13,403	6,319
Average Log-L	-1.1630	-1.2885	-1.2519	-1.2035	-1.2083

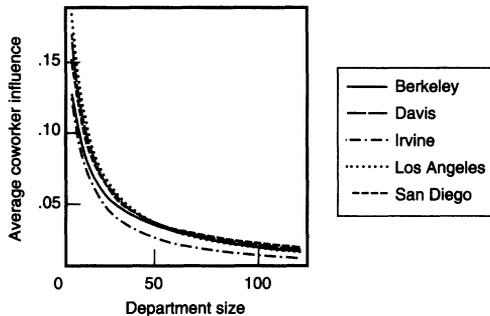
Note: Per-coworker effect is $\gamma_0 + \gamma_1 N + \gamma_2/N$. Standard errors in parentheses, adjusted for clustering by department. Coefficients on employee demographics omitted to save space.

Table 6 reports the average relative standard deviations of the different components in utility. (The standard deviations are taken across plans for each individual, and then the average of the ratio is taken across individuals.) With the exception of UCLA, the across-plan variation in social effects is typically half to three-quarters of the magnitude of utility variation based on individual characteristics. At UCLA, the variation in social factors typically exceeds any variation in individual characteristics. Compared to the standard deviation of the idiosyncratic preference term (normalized to $\pi/\sqrt{6}$ as usual), the standard deviation of social effects ranges between 14% and 28%.

What these numbers suggest is that while social effects play an important role, they are by no means the dominant determinants of health plan choice. With the exception of UCLA, individual characteristics (taken as a whole) appear to be more influential than social effects in new employees’ health plan decisions. Note that if the impact of social pressures on utility swamped any influence of individual characteristics, we would expect to see even more dramatic clustering than what we observe. In simulation tests (not reported here), I examined how large γ would have to be in order to cause widespread “tipping”—i.e., entire departments herding into the same plan—and found that γ must be at least 4 to generate such severe clustering.²²

□ **Social learning or common unobservables?** The signs and magnitudes of the estimates reported above seem quite reasonable when interpreted as reflections of social learning or peer

FIGURE 1
PER-COWORKER SOCIAL EFFECT AS FUNCTION OF DEPARTMENT SIZE



²² Plan choices were simulated for various values of γ while holding all other parameters fixed at their estimated values. These simulations generated choice patterns and clustering indexes similar to the observed data for values of γ below 3; for values above 5 virtually all employees within the same department would choose the same plan. Further details and a table describing these simulations are available from the author.

TABLE 6 Plan Enrollments by Campus (2000)

Campus	Average Ratio	
	$sd(\hat{\gamma} f)/sd(X\hat{\beta}_k)$	$sd(\hat{\gamma} f)/sd(\epsilon)$
Berkeley	.74	.24
Davis	.66	.23
Irvine	.57	.14
Los Angeles	1.28	.28
San Diego	.85	.23

Note: Based on estimates reported in Table 3.

influence. However, there are important reasons to be cautious about settling on such an interpretation. The models of the previous sections assumed away the “common unobservables” (η_{jk}) of equation (2), essentially ruling out the possibility that workers in the same department have similar unobserved preferences, so that their decisions will appear to be correlated even in the absence of any social learning. In the terminology of Manski (2000), the difficulty is distinguishing between “endogenous interactions” and “correlated effects.”

Several features of the data make it possible to address this difficulty directly. First, given the panel nature of the data, it is possible to leave the department-specific heterogeneity terms (the η_{jk} 's) in equation (2) and estimate them as fixed effects. The key assumption is that the department fixed effects are time invariant: although we allow for unobserved, department-specific heterogeneity in health plan preferences, we must assume that this heterogeneity does not vary with departments over time. Therefore, the estimate of the social-learning effect—i.e., the impact of coworkers' choices—is driven by changes from year to year in a department's collective health plan choices. So, for example, one employee may be hired into the department in 1997, when very few of her coworkers are in Pacificare. When another employee is hired in 1998 into the same department, she finds that a larger share of the employees are now in Pacificare (perhaps because some have switched, or simply because last year's cohort of new hires all selected Pacificare). The estimated social effect reflects how much more likely the latter employee is to choose Pacificare relative to the former employee.

Table 7 reports the relevant results from models with department-level fixed effects and “unweighted” social influence (the baseline model described earlier). Only new employees are used in the estimation, and the samples are limited to only those departments that hired at least one new employee in each of the six years in the sample period.²³ Based on the results in the table, it is clear that explicitly allowing for department-specific unobserved heterogeneity does not make the social effects go away. Although the estimated impact of coworkers' choices typically decreases once this heterogeneity is incorporated, the estimates of γ remain positive and statistically significant. The reduction in the estimated coefficient is modest for Irvine and San Diego, and for Davis the estimate actually increases. The similarity of these results to those reported in Table 3 supports the social-learning explanation over common unobservables.²⁴

In principle, the availability of data across different UC campuses permits another direct test of the common unobservables hypothesis. If the primary explanation for clustered choices within departments is department-specific heterogeneity—i.e., employees of the same department all “think alike” about health plan choices, so their decisions will appear to be correlated even in the absence of any social interaction—then we might expect one department's unobserved preferences to match those of the same department at a different campus. For example, if the

²³ Using additional departments that had fewer years with new hires is feasible, and the reported results are not sensitive to changing the requirement to four or more years (instead of six) with new hires. However, since the identification of the social effect comes from year-to-year variation in plan shares, using these departments adds many parameters to the model while adding relatively little identifying variation for the parameters of interest.

²⁴ The estimates of γ in the top panel of Table 7 differ from the numbers reported in Table 3 only because the samples are different.

TABLE 7 Estimates of Social Effects With and Without Department-Specific Heterogeneity

Plan	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Without dept. fixed effects					
Coworkers' share (γ)	2.056 (.203)	2.155 (.271)	1.247 (.198)	1.967 (.111)	2.120 (.213)
With dept. fixed effects					
Coworkers' share (γ)	1.072 (.338)	2.336 (.536)	.952 (.659)	1.070 (.196)	1.921 (.337)
<i>N</i>	4,940	5,009	3,013	11,286	5,105
Number of departments	97	74	50	152	80
LR test: no dept. effects	541.9 (.00)	522.4 (.00)	196.9 (.00)	987.7 (.00)	404.4 (.00)

Note: Models estimated using only those departments that hired new employees in all six years of sample period. Standard errors in parentheses, adjusted for clustering by department. Plan-specific coefficients on employee characteristics omitted to save space. For likelihood-ratio tests, χ^2 test statistics and *p*-values are shown. (The degrees of freedom differ depending on the number of department fixed effects estimated at each campus.)

unobserved preferences of physicists at UC San Diego lead them all to choose a particular plan, then physicists at UCLA should also be clustered on that same plan. In fact, this pattern is not observed in the data: same departments tend to be clustered on different plans at different campuses.²⁵ However, we should be skeptical of the premise of such a test: although it seems reasonable to expect employees of the same discipline to share similar unobservable characteristics across campuses, features of their choice sets are likely to vary across locations. So, for example, the fact that anthropologists at Berkeley predominantly chose Health Net while the anthropologists at UCLA predominantly chose UC Care may simply reflect that Health Net was a more attractive plan in northern California than in southern California. Although this is accounted for to some extent by the inclusion of campus-specific plan constants in the estimation, we should be reluctant to draw any conclusions from these cross-campus tests.

As has been noted by previous authors,²⁶ the potential confounding influence of common unobservables suggests looking for patterns within subgroups of the main social unit. If we have prior information that social interaction takes place primarily within certain subgroups, then we should expect behavior to be most correlated within those subgroups—i.e., individuals would be influenced by the decisions of peers in the same subgroup, but not by the decisions of peers in other subgroups. In the context of universities, the most natural social subnetworks to consider are faculty and staff within academic departments: to the extent that faculty socialize primarily with other faculty, we expect a new faculty member's health plan choice to be influenced most by other members of the faculty in her department. Table 8 reports estimates of the baseline model allowing the social effect to vary across subgroups:

$$u_{ik} = x_i' \beta_k + \left(\sum_{j=1}^{n_i} \gamma_{g(i),g(j)} d_{jk} \right) + \varepsilon_{ik}, \tag{4}$$

where *j* indexes the *n_i* coworkers of employee *i*, *d_{jk}* is an indicator equal to one if employee *j* was enrolled in plan *k*, and *g*(·) is either faculty (*f*) or staff (*s*). So, for example, γ_{ff} represents the influence of a faculty coworker's health plan choice on a new faculty employee's choice, while γ_{fs} represents the influence of a staff coworker on a faculty member's choice, and so on. The

²⁵ Discrete choice models were estimated using 21 academic departments that exist at three or more of the campuses, first using (department × campus) fixed effects and then using department fixed effects—i.e., imposing the restriction that unobserved preferences are identical for same departments at different campuses. Likelihood ratio tests soundly reject the validity of the restriction ($\chi^2_{158} = 310.04, p = .000$).

²⁶ See, for example, Munshi and Myaux (2000) and Dufo and Saez (2002).

TABLE 8 Academic Departments Divided into Faculty and Staff Subgroups

Plan	Campus				
	Berkeley	Davis	Irvine	Los Angeles	San Diego
Own-group:					
Faculty-Faculty (γ_{ff})	1.316 (.389)	.588 (.263)	1.403 (.318)	1.256 (.213)	2.086 (.370)
Staff-Staff (γ_{ss})	1.473 (.183)	1.816 (.193)	1.282 (.167)	1.804 (.096)	1.953 (.165)
Cross-group:					
Faculty-Staff (γ_{fs})	.006 (.373)	.436 (.360)	.548 (.347)	1.018 (.365)	.786 (.361)
Staff-Faculty (γ_{sf})	.275 (.141)	.219 (.147)	.188 (.135)	.175 (.070)	.266 (.118)
<i>N</i>	6,345	6,911	3,780	13,403	6,319
Average log-likelihood	-1.1634	-1.2896	-1.2477	-1.2015	-1.2081
LR test: ($\gamma_{ff} = \gamma_{ss}$) & ($\gamma_{fs} = \gamma_{sf}$)	1.33 (.52)	16.86 (.00)	1.62 (.44)	1.31 (.01)	2.10 (.35)
LR test: ($\gamma_{ff} = \gamma_{fs}$) & ($\gamma_{ss} = \gamma_{sf}$)	5.3 (.00)	85.4 (.00)	3.3 (.00)	346.3 (.00)	113.6 (.00)

Notes: Standard errors in parentheses, adjusted for clustering by department. Plan-specific coefficients on employee characteristics omitted to save space. "Faculty-Faculty" reflects impact of faculty coworkers' choices on new faculty's choices. "Faculty-Staff" indicates impact of staff coworkers' choices on new faculty's choice. For likelihood-ratio tests, χ^2 test statistics and *p*-values shown.

table indicates that own-group (either faculty-to-faculty or staff-to-staff) effects are positive and statistically significant, while the cross-group effects are typically smaller and/or indistinguishable from zero—a pattern that is fully consistent with the social-learning hypothesis but inconsistent with department-specific common unobservables. However, as was noted previously, unobserved heterogeneity at the subgroup level would still be consistent with the observed patterns. That is, it could just be that physicists all think alike about health plans, but they don't necessarily think like physics department staff. On this point, it is somewhat informative that the measured faculty-to-faculty influence is not larger than the staff-to-staff influence. One could argue that since faculty members in the same academic department usually have similar training and background, they are more prone to think alike than staff members (whose training and background are typically more varied). If so, the common unobservables problem would be most relevant for faculty. However, at the only two campuses with a statistically significant difference between γ_{ff} and γ_{ss} , the faculty-to-faculty effect is actually smaller than the staff-to-staff effect.

As a final reality check, we can estimate the baseline model reported in Table 3 with an intentionally misspecified social group. That is, instead of asking whether an individual's plan choice is influenced by coworkers within her department, we can ask whether her choice was influenced by coworkers from another department. Not surprisingly, if departments are paired randomly, the estimated impact of pseudo-coworkers' choices is essentially zero. (The point estimates of γ are small and statistically insignificant.) However, if departments are paired with departments that are similar—based on whether they are academic versus nonacademic departments, and on department size and average salary—then the estimated influence of pseudo-coworkers' choices is positive and statistically significant, though considerably smaller than the estimates reported in Table 3 or in Table 7. (The point estimates of γ are around .6 for each campus.) This implies that while social effects are clearly present, selection on unobservables is not entirely irrelevant, which reinforces the importance of controlling for those unobservables directly (as in the fixed-effects models reported in Table 7).

4. Discussion and conclusions

■ 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 appear

to depend on peers' information and observed actions, the opposite may be true of employees who switch plans. In particular, switches may reflect decisions made based on new private information—e.g., individuals may “go with the crowd” until they obtain enough information to determine their individual optimum. Although it is beyond the scope of the present article to analyze this hypothesis in detail, a few basic facts from the panel data are suggestive. First, the patterns described above indicate that new hires are disproportionately likely to choose plans that are “dominant” in the departments where they are hired (dominant being defined as “most overrepresented relative to what we’d expect based on employee characteristics”). Second, the most common time to observe a switch in health plan is in an employee’s second or third year of employment. Finally, for new hires who chose plans that were popular in their departments, second- or third-year switches are roughly twice as likely when the plan initially chosen appears to be a bad match based on the hired employee’s characteristics (relative to the case when the popular plan also happened to be a good match).

Such patterns raise an important question about this study’s findings: Is the social learning beneficial? Information sharing can clearly increase welfare when the shared information enables individuals to make better decisions. On the other hand, inefficient herding is also possible: social learning can lead groups to coordinate on suboptimal choices. The data here are insufficient to conclusively resolve this question. Given that the measured impact of peers’ decisions is moderate, the potential for inefficient herding is probably small. However, if the social effects reflected only efficient information sharing, we might expect individuals with similar characteristics to ultimately converge on the same plan—but this is not what we observe in the data.

This article confirms the importance of social effects in a particular context: individuals’ choices of employer-sponsored health plans. The quantitative results indicate that decision makers are influenced substantially by their coworkers’ choices, but that these effects do not necessarily dominate other factors that are relevant to the decision (such as age, income, or family status). Though the quantitative estimates are specific to this context, the findings are illustrative of the role of social learning in markets for experience goods more generally—and in many markets one can imagine that social effects may be much more salient.

The implications of social learning for firm behavior are manifold. For example, social effects may dramatically increase firms’ incentives to promote their products, since (roughly speaking) by gaining one customer a firm can gain more than one in expectation. What may be most important is that individuals’ tendency to follow the crowd can potentially give considerable market power to firms that obtain dominant market shares. In the present setting, if virtually all of the members of a department have chosen the same plan, then the social influence may overpower any individual incentives to switch plans when the provider raises the premium. In this respect, social-learning effects closely resemble network effects, and from the firm’s perspective the implications are similar. Previous theoretical research about network effects and their impact on firms’ product-market behavior (e.g., Katz and Shapiro, 1985) may therefore be applicable to markets with social learning. However, some patterns of interaction described in this article seem to be unique to social-learning models. Understanding the implications of such patterns for supply-side behavior would be a valuable objective for future research.

References

- BANERJEE, A.V. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, Vol. 107 (1992), pp. 797–817.
- BAYER, P., ROSS, S., AND TOPA, G. “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes.” NBER Working Paper no. 11019, 2005.
- BERTRAND, M., LUTTMER, E., AND MULLAINATHAN, S. “Network Effects and Welfare Cultures.” *Quarterly Journal of Economics*, Vol. 115 (2000), pp. 1019–1055.
- BIKHCHANDANI, S., HIRSHLEIFER, D., AND WELCH, I. “A Theory of Fads, Fashions, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy*, Vol. 100 (1992), pp. 992–1026.
- BROCK, W.A. AND DURLAUF, S.N. “Discrete Choice with Social Interactions.” *Review of Economic Studies*, Vol. 68 (2001), pp. 235–260.
- BUCHMUELLER, T.C. AND FELDSTEIN, P.J. “The Effect of Price on Switching Among Health Plans.” *Journal of Health Economics*, Vol. 16 (1997), pp. 231–247.

- CASE, A. AND KATZ, L. "The Company You Keep: The Effect of Family and Neighborhood on Disadvantaged Youths." NBER Working Paper no. 3705, 1991.
- CONLEY, T.G. AND UDRY, C.R. "Learning About a New Technology: Pineapple in Ghana." Mimeo, Yale University, 2000.
- DUFLO, E. AND SAEZ, E. "Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues' Choices." *Journal of Public Economics*, Vol. 85 (2002), pp. 121–148.
- ELLISON, G. AND FUDENBERG, D. "Rules of Thumb for Social Learning." *Journal of Political Economy*, Vol. 101 (1993), pp. 612–643.
- AND GLAESER, E.L. "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy*, Vol. 105 (1997), pp. 889–927.
- FARRELL, J. AND SALONER, G. "Standardization, Compatibility, and Innovation." *RAND Journal of Economics*, Vol. 16 (1985), pp. 70–83.
- GLAESER, E.L., SACERDOTE, B., AND SCHEINKMAN, J.A. "Crime and Social Interactions." *Quarterly Journal of Economics*, Vol. 111 (1996), pp. 507–548.
- HONG, H., KUBIK, J., AND STEIN, J. "Social Interaction and Stock Market Participation." NBER Working Paper no. W8358, 2001.
- KATZ, M. AND SHAPIRO, C. "Network Externalities, Competition, and Compatibility." *American Economic Review*, Vol. 75 (1985), pp. 424–440.
- MANSKI, C.F. "Economic Analysis of Social Interactions." *Journal of Economic Perspectives*, Vol. 14 (2000), pp. 115–136.
- MUNSHI, K. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." Mimeo, University of Pennsylvania, 2000.
- AND MYAUX, J. "Social Change and Individual Decisions: With an Application to the Demographic Transition." Mimeo, University of Pennsylvania, 2000.
- TAMER, E.T. "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria." Mimeo, Princeton University, 2000.
- WOITTEZ, I. AND KAPTEYN, A. "Social Interactions and Habit Formation in a Model of Female Labour Supply." *Journal of Public Economics*, Vol. 70 (1998), pp. 185–205.
- YOUNG, H.P. AND BURKE, M.A. "Competition and Custom in Economic Contracts: A Case Study of Illinois Agriculture." *American Economic Review*, Vol. 91 (2001), pp. 559–573.