

Participation and Investment Decisions in a Retirement Plan: The Influence of Colleagues' Choices

Esther Duflo, MIT and NBER

and

Emmanuel Saez, Harvard University and NBER*

May 25, 2000

Abstract

This paper investigates whether peer effects play an important role in retirement savings decisions. We use individual data from the staff of a university to study whether individual decisions to enroll in a Tax Deferred Account plan sponsored by the university (and the choice of the mutual fund vendor for people who choose to enroll) are affected by the decisions of other employees in the same department. To overcome the identification problems, we separate the departments into sub-groups (along gender, status, age, and tenure lines) and we instrument the average participation of each peer group by the salary or tenure structure in this group. Our results suggest that peer effects are important. We find significant own-group peer effect on participation and on vendor's choice, but no cross-group peer effects. (*JEL* D83, I22)

*MIT, Department of Economics E52-252g, 50 Memorial Drive, Cambridge, MA 02142, eduflo@mit.edu. Harvard University, Department of Economics, Littauer Center, Cambridge, MA 02138, saez@fas.harvard.edu. We thank Abhijit Banerjee, David Cutler, Jonathan Gruber, Lawrence Katz, Michael Kremer, Philip Lima, Brigitte Madrian, Sendhil Mullainathan, Kaivan Munshi, Linda Nolan, James Poterba, and Rosemary Rudnicki for helpful discussions, and Kirsten Carter, Regina Perris, and Polly Price for giving us access to the data.

1 Introduction

Low levels of savings in the United States have generated substantial interest in the question of what determines savings decisions. A vast literature has studied the impact of Tax Deferred Accounts (hereafter, TDA), such as Individual Retirement Accounts (IRA) and 401(k)s, on retirement savings decisions,¹ and, concurrently, the impact of these plans' features on enrollment and contribution rates. A number of studies attempted to assess the effect of economic incentives on individual behavior and found mixed evidence. The presence of a matching contribution from the employer has generally been found to be correlated with higher participation rates, but the level of the match rate does not seem to matter.² As Bernheim (1999) points out, matching also serves as a device to focus the employees' attention. This suggests that pure economic incentives are not sufficient to explain savings behavior. Recent studies emphasize the role of non-economic incentives such as financial education and inertia. Madrian and Shea (2000b) show that default rules have an enormous impact on employees' participation, contribution, and asset allocation. When they are enrolled by default in a TDA, very few employees opt out. Further, most employees do not change the default contribution rate or the default allocation of assets. Bernheim and Garrett (1996) and Bayer, Bernheim and Scholz (1998) study the role of financial education. They present evidence that financial education tends to be remedial³ but that it increases participation in the plan, suggesting that employees may not be able to gather the necessary information on their own.

This paper contributes to this literature by studying the role of peer effects in TDA participation and decisions related to the plan. There has never been a study of peer effects on saving decisions. This is surprising, because the theoretical literature suggests at least two reasons why peers play a role in this context. First, the plans are sufficiently subtle that their advantages are not obvious to someone who has not thought carefully about it. Even when people choose to participate, they may lack the information necessary to make investment decisions. The evidence presented by Madrian and Shea (2000b) suggests that a large proportion of people do not

¹See Poterba, Venti and Wise (1996) and Engen, Gale and Scholz (1996) in a special issue of the *Journal of Economic Perspectives*.

²See, e.g. Papke (1994), Papke, Petersen and Poterba (1993), and Kusko, Poterba and Wilcox (1994).

³Employers resort to it when they fail discrimination testing because the contribution rates of the not highly compensated employees are too low.

think about these decisions at all. The literature on informational cascades (Banerjee (1992), Ellison and Fudenberg (1993)) provide reasons why information (correct or not) obtained from co-workers may be an important factor in deciding whether to participate and how to invest—giving rise to peer effects. Second, savings decisions may be influenced by social norms or beliefs about social norms. By observing co-workers, people can learn about the proper behavior of their social group, as emphasized by models of conformity (e.g. Bernheim (1994)): Individuals may want to maintain the same consumption level as what is common in their social group.

There is a growing empirical literature on peer effects which essentially focuses on social behavior, and the adoption of new technologies.⁴ Manski (1993) provides a formal exposition of the econometric issues that identifying peer effects involves. Correlation of behavior within peer groups is not necessarily due to the fact that members of the group directly influence each other. First, members of the same group share a common environment, which may influence their behavior. Second, except when individuals are randomly assigned to a peer group, people with similar preferences tend to belong to the same group. Both of these generate a correlation between group behavior and individual behavior which does not indicate any causal relationship between the two. Finally, there may be a causal relationship between the characteristics of the peer group members and individual behavior which does not reflect either learning or conformity. For example, redlining by areas means that the racial composition of one's neighborhood directly affects the probability that one can get a mortgage. Similarly, employees working in firms where other people are well paid may directly benefit from some of these advantages. This is what Manski (1993) calls an exogenous social effect.

In this paper, we ask whether the decisions of employees of a large university to enroll in the TDA plan, and the vendor they choose once enrolled, are affected by the decisions of their colleagues in the same department. We begin by presenting an intriguing example, namely the differences in participation rates among the university's libraries. Although average staff salary and experience are very similar across libraries, participation rates are very different. This suggests that the correlation of employees behavior within libraries may be due to peer effects.

⁴See for example, Case and Katz (1991) and Evans, Oates and Schwab (1992) on teenagers' behavior, Sacerdote (2000) on college students behavior and choices, Bertrand, Mullainathan and Luttmer (1998) on welfare participation, Munshi (2000b) on contraception, and Besley and Case (1994), Foster and Rosenzweig (1995) and Munshi (2000a) on technology adoption in developing countries.

In the remainder of this paper, we focus on the decisions of the administrative and support staff of the university as a whole. There are several reasons why the identification of peer effects is easier in this context than in other situations previously studied. First, the employees share a common program, centrally organized by the university. Benefits informational sessions are identical for all departments in the university. The particular department in which one works therefore does not affect the level of inputs provided by the firm to help the employees make their TDA decisions. Second, employees do not choose to work for a particular department because it made enrollment in the TDA plan easier. It is still possible for propensity to save to be correlated within departments. For example, economists probably know more about TDA plans than physicists, and thus are more likely to participate even if we control for earnings levels. Even when we restrict our sample to the staff, we may not remove all of this correlation. Third, once we control for individual wages or tenure, the average wage or tenure in the department does not directly affect individual enrollment decisions. We follow Case and Katz (1991), and use this assumption to construct instruments for the average participation in the plan. The instruments can still be invalid if there is a correlation between average wage (or tenure) in a department and the individual's unobserved propensity to save even after controlling for individual wage and tenure.

Fourth, presumably, individuals interact mostly with co-workers who share observable characteristics such as gender, age, or tenure. Put another way, women are more likely to talk to women, men to men, and newly hired employees to newly hired employees. Therefore, it is plausible that the relevant peer group of an individual is a sub-group of his department. We use this presumption to construct another test of the identification assumption. We regress individual participation on average participation in his own sub-group and the average participation in the other sub-groups. If there is a correlation between the instruments and the error term at the department level, we would see a (spurious) positive coefficient for the average decision of the other sub-group in the department.

Lastly, we study the choice of the mutual fund vendor in addition to the participation decision. Because vendors offer similar services at similar costs, we might think that employees do not feel very strongly about any one vendor, and that if some have a preference for one vendor over another, these preferences are probably not correlated within departments. If, using the

aforementioned techniques, we find a positive association between the choice of vendors within sub-groups and departments, it should reinforce our confidence in the previous findings.

The remainder of the paper is organized as follows. In section 2, we provide evidence from the university’s libraries as an introductory example. Section 3 describes the features of the university’s TDA plan and the data. In section 4, we present the results on the participation decisions. In section 5, we turn to the choice of vendor. We find strong evidence of peer effects for both participation and vendor choice. Section 6 concludes.

2 Case Study: Libraries

In table 1 we present some preliminary but suggestive evidence. It presents the contribution rates, salary, and tenure of the staff in the university’s 11 independent libraries that are jointly administered by a central library administration.⁵ Libraries differ in the number of staff members, but the composition of the staff is similar across libraries. Salary and years of services are very similar in all libraries. There are reasons to believe that the assignment of employees to a particular library is not correlated with unobserved characteristics related to savings. Most library jobs currently listed in the university’s employment web site do not require a specific competence related to a domain, but a general education and library experience. The central library administration hires new library employees. Initial assignment to a library, therefore, is mostly determined by the opening of a position suitable to the applicant at the time he or she applied. Moreover, year to year transitions from one library to another are extremely rare. From a total of 1,800 observations (around 450 employees in libraries observed four times over two years), there is only one occurrence of an employee switching from one library to another. In summary, there is no *a priori* reason to think that the library employees’ propensity to save should be correlated across libraries. In the absence of peer effects, we should expect relatively similar participation rates from one library to another.

However, participation rates differ substantially from one library to another. The rates vary from 0.14 to 0.73. Is there too much variance in the distribution of participation rates across libraries relative to what we would expect in the absence of correlation of behavior

⁵There are other libraries in the university that are administratively attached to a specific department, and are not part of the central library administration.

within libraries?⁶ Under the null hypothesis that the individual probabilities of participation are independent and given by the empirical average participation rate p over all libraries, the variance of participation rates across libraries would be equal to 0.239. However, the actual empirical variance is 0.418.⁷ Using a simple bootstrap method, we find that the null hypothesis is rejected with a p-value of 0.049. Alternatively, an OLS regression of individual participation on the average participation of other employees in the library leads to a coefficient of average participation (0.31) that is significant at the 5% level, and is not affected by controlling for wages or years of service.

Simple evidence from a very small subsample (one wave of data and less than 5% of the university's employees) suggests that peer effects are present. Of course, there are other possible interpretations of this result. First, some libraries are more prestigious than others, and therefore the human resources department of the library administration may direct the most competent applicants towards those libraries (the fact that salaries and tenure do not vary much from one library to another is comforting but the staff composition may still differ along unobserved skill dimensions correlated to savings). Second, some library staff *do* have special skills or characteristics (for example, employees at the Oriental Studies Library are more likely to be Asian-Americans), which may be correlated with their propensity to save. Nevertheless, this dispersion is striking.

In the following sections, we present evidence on the importance of peer effects on the decision to enroll in the TDA and on the choice of vendor using data from the university as a whole, where the assumption of random assignment to a department cannot be made.

⁶Similarly, Glaeser, Sacerdote and Scheinkman (1996) interpret the excessive variance in crime rates across cities as evidence of peer effects.

⁷Under the null hypothesis, the variance of the average participation, P_x , in a department with N_x employees is $p(1-p)/N_x$. Therefore, the variance of $\sqrt{N_x}P_x$ is $p(1-p)$ for all x . Empirically $p(1-p) = 0.239$ and the variance of $\sqrt{N_x}P_x$ across x is 0.418.

3 Description of the TDA Plan and the Data

3.1 Features of the TDA Plan

The university we study has approximately 12,500 employees. Roughly 20% of the employees are members of the faculty. The university provides retirement benefits to its employees through a traditional pension plan and a complementary Tax-Deferred-Account (TDA) plan.

Part of the traditional pension plan is a Defined-Contribution (DC) plan where a given percentage of an employee's salary is put into an individual mutual fund account run by the fund manager.⁸ The firm contributes 3.5% of salaries into this DC plan for staff employees and 5 to 10% (depending on tenure and age) for faculty. There is a one year waiting period for the DC plan benefit to begin.

Employees can also contribute to a TDA plan, a 403(b) plan⁹ which has no waiting period. Every employee can contribute to the 403(b) plan any percentage of their salary up to the IRS limit (approximately \$10,000 per year for each individual). The university does not match contributions.

In both the DC and the TDA plans, employees can choose where to invest their contributions from any number of four different vendors. Each vendor provides a large selection of mutual funds (around 25 each) that include money-market funds, bonds, and stocks (both U.S. and foreign). Each of the four vendors offer very similar services. All vendors allow customers to change their portfolios in a very flexible way through the phone or the internet. We will concentrate on the three biggest vendors which attract over 90% of total contributions.¹⁰ We denote these three vendors R, D, and V.

3.2 Summary Statistics

The university provided us with individual data on TDA participation and contributions. The university collected four waves of data (October 1997, June 1998, October 1998, and June 1999)

⁸Staff employees have an additional Defined Benefits plan in addition to the DC plan.

⁹403(b) plans are very similar to the better known 401(k) plans but their use is restricted to not-for-profits firms.

¹⁰The fourth vendor represents only 6% of the contributions. The remaining 4% are spread among a few vendors which are no longer offered.

on all the employees. Individual identifiers are provided so that the four waves can be linked. We use the four waves together and correct the standard errors for clustering at the individual level.

A number of variables are included for each employee and each wave. Table 2 provides summary statistics of the variables we use in this study. It displays means and standard deviations for three groups of employees. Column (1) is the complete sample (both staff and faculty). We exclude from this sample all employees working in food and custodial services, because their wages and contribution levels are substantially below those in the other departments. We exclude the business school as well because we did not obtain the breakdown by departments within the school.¹¹ In total, these excluded observations represent slightly less than 10% of the initial sample. Because there are strong reasons to think that faculty are sorted into departments in a way that is correlated to their propensity to save, we restrict the analysis of peer effects for the participation decision to the staff. Column (2) presents descriptive statistics for this sample. Since we are restricted to individuals who participate in the TDA when examining the choice of vendor, we will present results for all staff members who participate. Finally, we present the descriptive statistics for faculty in column (3).

Panel A, table 2 displays demographic and compensation characteristics of the university's employees. The average salary among staff is a little over \$39,000. The percentage of male employees in our main sample is 39%. The average age is 40 and the average tenure is 7.3 years. Panel B, table 2 presents information on TDA plan participation and choice of vendors. The average participation rate is 33% among staff. The share of contributions in vendors D, R, and V are 34%, 34%, and 22%, respectively. Panel C, table 2 presents information on departments. Our sample is divided into 420 departments. The average number of employees per department is 32, and the median is 12. The smallest department has 2 members.

4 Peer Effects in Participation Decisions

In this section, we first show that participation rates are correlated across departments, and we then present evidence which suggests this correlation is, at least in part, driven by social effects.

¹¹As a result, the number of employees in the business school is over 700, much more than the next largest department.

4.1 OLS Results

Each individual in the firm is characterized by a vector (y, x, Z, u) . y is the outcome of interest. In this paper we consider two potential outcomes of interest. First, we study participation decisions. In that case, y is a dummy for participation in the TDA. Second, we study the choice of the mutual fund vendor conditional on participation. x is the department where the employee is located. (Z, u) are individual characteristics that affect the outcome y . The characteristics Z are observables (salary, gender, age, and years of service). u is an unobservable scalar which represents unobservable characteristics that might affect the outcome y , such as propensity to save, that are not captured by the observables. We assume a linear specification and we seek to estimate the model,

$$y_i = \alpha + \beta \hat{E}_{-i}(y|x) + Z_i \eta + u_i, \quad (1)$$

where i is an individual observation, and

$$\hat{E}_{-i}(y|x) = \sum_{j \in x \setminus \{i\}} y_j / (N_x - 1)$$

is the average of y in department x (excluding individual i). N_x denotes the number of individuals in department x . Excluding the observation of individual i when computing the expectation avoids a purely mechanical correlation between average departmental participation and individual participation.

The results of estimating this equation by OLS are presented in column (1), table 3. We control for gender, dummies for each age decade, tenure dummies, and salary dummies indicating in which decile of the university-wide distribution of salaries the individual falls. The OLS coefficient, in Panel A for the staff sample, shows that, controlling for age, tenure, and salary, there is a strong correlation between individual participation decisions and average participation in the department. Each additional percentage point of participation in the department is associated with a 0.27 percentage point increase in the individual's probability of participating in the TDA. In table A1, we present the coefficients of the control variables.¹²

¹²Women tend to contribute more than men. Participation increases with age, tenure (especially during the first 4 years), and with wages.

The OLS coefficient is biased and cannot be interpreted as evidence of the presence of peer effects when the error term u is correlated with the variable of interest $\hat{E}_{-i}(y|x)$. This correlation is present when some of the unobservable characteristics which influence an individual's participation (or vendor's choice) are correlated across departments even in the absence of any peer effects. Following Manski's (1993) terminology, these are called 'correlated effects'. For example, there could be direct sorting on the propensity to save. Alternatively, most departments are responsible for hiring their staff members. Therefore, employees may all be similar because they have been chosen by the same person, and each person in charge of hiring may emphasize different forms of competence, some of which are quite possibly correlated with propensity to save (for example 'good' departments may hire 'competent' people, and competent people may make good financial decisions).

4.2 Two-Stage Least Squares

To rule out the possibility that the correlation of behavior within departments is driven entirely by unobserved characteristics, the ideal experiment would be to allocate employees randomly to departments (an example of random allocation of roommates is studied in Sacerdote (2000)), or to induce a modification of the contribution rate of a random subset of employees in some departments. One would then need to compare the participation of the non affected employees between those departments and the departments where no intervention took place.¹³

In the absence of such an experiment, average exogenous characteristics of the groups could be used as instruments for the average participation (for a previous application of this strategy, see Case and Katz (1991)). As we saw in the previous subsection, tenure and wages are strong determinants of participation in the plan. Therefore, average wages and or tenure in the department are also strongly correlated with average participation. Thus $E(Z_1|x)$, where Z_1 is a subset of the characteristics of Z such as wages or tenure, can potentially be used as an instrument. The F-statistic of the first stage correlation between $\hat{E}_{-i}(y|x)$ and $E(Z_1|x)$ is strong. These variables are valid instruments if they do not directly affect the participation rates (through what Manski calls 'exogenous' social effects), and if they are not correlated with the unobserved determinant of savings, two points to which we will return below.

¹³We will describe such an experiment in more detail in the conclusion.

In columns (2), (3), and (4) of table 3 we re-estimate equation (1), using, respectively, the proportion of individuals in the department who fall in any given decile of the university-wide distribution, the proportion of individuals in each tenure category, and both together, as instruments for average participation. The three coefficients obtained by instrumental variables are similar, and they drop from 0.27 to between 0.13 and 0.15, which indicates that the OLS coefficient is upward biased, probably due to omitted correlated effects. The coefficient remains sizeable and significant when the two instruments are used together. The effect on participation of raising the average participation in the department by one percentage point is larger than the effect of moving from the first to the fourth decile in the wage distribution. The similarity of the estimates obtained with the two alternative instruments is reassuring: An identification test would not reject the joint validity of the instruments. At the bottom of table A1, the F-statistics of the first-stage are reported. These statistics are large (between 15 and 30) showing that the instruments are significantly correlated with the participation rate in each department.¹⁴

These 2SLS results alone do not constitute definitive evidence, because the two conditions necessary for the validity of the wage and tenure variables as instruments may fail. First, there may be a direct exogenous effect of average wages or tenure in the department on an individual's participation, even after conditioning for one's wage. For example, high wage employees can require their administrator to collect and disseminate information on available benefits, which may trickle down to low-paid employees as well. Second, unobserved characteristics correlated with propensity to save (e.g. "competence") could well be correlated with the department's average wage, even after conditioning for an individual's wage. For example, professors' salaries in the well renowned departments may be higher than in other departments, and these departments may also be able to hire more competent staff.

A simple way to assess whether 2SLS helps to solve the problem is to replace the participation variable with age variables. Obviously, age cannot be affected by the age of colleagues through peer effects. An OLS regression of individual age on average age in the department (controlling for gender and salary) produces a positive and significant coefficient. However, when average age is instrumented with average salary in the department, the coefficient is much smaller and

¹⁴The F-statistics are obtained by running the first stage at the department level. They are, therefore, conservative.

insignificant. This suggests that the IV strategy can be successful in removing correlated effects. In the next two sections, we provide additional evidence to reinforce our confidence in these results.

4.3 Differential Effects According to Department Size

Our main concern is that the OLS and 2SLS results may be picking up a correlation between unobserved tastes or unobserved inputs within departments, rather than direct influence from colleagues. In other words, average participation in the department (or the average wages and tenure that we use as instruments) might be a proxy for a department-level fixed effect. How can we discriminate between the two hypotheses (the average participation in the department is a proxy for a departmental fixed effect versus it reflecting the influence of others on an individual's behavior)?

We cannot provide a definitive test based on our observational data, but the following observation suggests we should consider how our estimates vary with the size of the department.¹⁵ Let us consider the OLS case. Note that under both hypotheses, the average participation is a proxy (measured with error) for an underlying latent variable (the behavior of the actual peers of an individual in one case, and the common propensity to save in the department in the other case). Therefore, the OLS regression in both cases gives a downward biased estimate of the parameter of interest. However, the two hypotheses have symmetrical implications for the relationship between department size and the size of the OLS coefficient. In the case where the correlation is due entirely to the presence of a departmental fixed effect, we should see the coefficient being closer to one in larger departments, because the department's underlying propensity to save will be better measured in larger departments. In the case where the coefficient picks up a causal effect of one's colleagues' behavior, we are likely to see the relationship go the other way: The *actual* peer group of an employee is presumably only a subset of all the colleagues in his or her department, more or less fixed in size. So as departments get larger, the measurement error increases, and the attenuation in OLS should become larger as well. Therefore, if the correlation is largely spurious, the coefficient should be bigger in large departments, whereas if it is due to social effects, the coefficient should be smaller in large departments. The same argument can

¹⁵We are grateful to Sendhil Mullainathan for pointing us in this direction.

be applied to the 2SLS results.¹⁶

In panels B and C of table 3, we present the results for the departments with less than 20 employees and more than 20 employees, respectively.¹⁷ The results in column (1) confirm that the OLS coefficients presented in panel A reflect something other than peer effects: The coefficient is larger in large departments than in small departments. However, the 2SLS results are more encouraging: The estimated effect is twice as large in small departments than in large departments. This can be seen as suggestive evidence that 2SLS estimates at least partially overcome the correlated effects problem.

4.4 Looking at Sub-Groups Within Departments

Actual peer groups are in many cases smaller than departments. As noted by Munshi (2000b), this fact can be used to help identify peer effects. If peer groups are only a subset of each department, there is an a priori restriction on the pattern of peer effects: There should be no effects of the participation of members of the other sub-groups on the members of one sub-group. There are sub-groups within departments where peer effects can be expected to be stronger than for the department as a whole. For example, newly hired employees may talk more to other newly hired employees than to established employees, and vice-versa. Moreover, if established employees have already made their decisions, they are not likely to be affected by what newly hired employees do (since decisions are rarely reversed after a length of tenure). Women probably talk more to women than to men, and men more to men than to women.

Here, we follow an idea developed by Munshi (2000b) which proposes to regress individual participation for each sub-group separately on the participation in their own and in the other sub-group.¹⁸ If there is a department-level correlated effect, it should cause the coefficient of average cross-group participation to be positive (even in the absence of peer effects).¹⁹ In other words we run,

¹⁶In the case of 2SLS, attenuation bias is present because both the instruments and the average participation are measured with error and these errors are correlated.

¹⁷The median is 12. We obtain similar results using other cut-off size levels.

¹⁸Bertrand et al. (1998) exploits a related idea: They study whether the number of welfare participants who speak an individual's language affect his/her participation, after controlling for the fraction of participation in his/her area of residence.

¹⁹This is formally proven in an earlier version of this paper.

$$y^k = \beta^k E(y|x, k) + \gamma^k E(y|x, \bar{k}) + Z\eta^k + u^k, \quad (2)$$

where k is the sub-group within a department (we assume that each department is partitioned into two sub-groups with $k = 0$ or $k = 1$) and y^k is the outcome of an individual in group k . We note \bar{k} the complement of k . As before, we allow for the possibility of a correlation between u^k and x . If we believe that cross-group effects are zero (that is, the parameter γ^k is equal to 0) and that the error terms in each sub-group are correlated then we can estimate equation (2) by OLS or 2SLS, and test whether the estimate $\hat{\gamma}^k$ is 0.²⁰ If $\hat{\gamma}^k$ is different from 0, it will indicate that there are correlated or exogenous effects at the departmental level, and this biased our previous estimate. If $\hat{\gamma}^k$ is equal to 0, it will indicate that previous estimates were not biased by any correlated or exogenous effects that are at least in part *common* to the entire department.

There could still be a problem with the OLS version of this test: If women and men are doing different jobs, a different person could be in charge of hiring them or the same person could emphasize different skills (men and women would then in effect form two different ‘departments’). The same could apply to tenure, if the person in charge of hiring had changed (although it is difficult to imagine that it would happen in all departments at the same time). However, when we combine the sub-groups and the 2SLS strategy, we would need to tell complicated stories to explain why all cross-group effects are zero. For example, when we use the salary instrument, we allow for the fact that a woman with high propensity to save would be more likely to work in a department where salaries are high. For example, if highly paid faculty hired high savings employees, it would lead to a positive $\hat{\gamma}^k$, unless only *women* faculty were involved in the hiring of women employees, an implausible assumption. For all the sub-groups we consider, it seems reasonable to consider that the error term should be correlated across sub-groups within department. Taking the results together should therefore give us a good idea of the presence of peer effects in retirement plan decisions.

In table 4, we present the results on peer effects among sub-groups. The first two columns present the results for the participation of individuals in group 1, and the last two columns

²⁰The instruments are constructed as the expectation of the subset Z_1 in Z in each sub-group for each department ($E(Z_1|x, k)$).

present the results for the participation of individuals in group 2. In each panel, the first line displays the effect of the average participation of members of group 1 and the second line displays the effect of the average participation of members of group 2. The third line reports the p-value for the test that the two coefficients are equal. Columns (1) and (3) report the OLS results, columns (2) and (4) report the 2SLS results using both salary and tenure as instruments.

In panel A, we present the effect of the average participation broken down by gender. In columns (1) and (2) of table 3, we see that participation of the female employees is significantly affected by the average participation of other women (the OLS and IV coefficients on female participation are respectively 0.20 and 0.32 and all are significant) but not by that of men (the coefficients are respectively 0.07 and -0.115 and none of them are significant). The equality of coefficients of own- and cross-groups is rejected in both cases. Symmetrically, the participation of the male employees seems to be affected by participation of other men, but not by the participation of women, although the OLS coefficients are not statistically different from each other.

In panels B, C, and D, we repeat the exercise by breaking the sample according to tenure (below 3 years and above 7 years), age (below or above 36), and faculty status (we still focus on the decisions of the staff). In all cases, we might expect to find positive cross-group effects, because there is no strong a priori reason to believe that peer effects are completely absent across these sub-groups. In all cases, however, cross-group coefficients are very small (sometimes negative) and insignificant, while own-group effects are always positive. Equality of the coefficients is rejected at the 10% level or best in most cases. When equality is not rejected, it is because the coefficient of the own-group is smaller, and not significantly different from zero, never because that of the cross-group effect is large.

These results, taken together, suggest both that there are no peer effects across these sub-groups within the departments and that the 2SLS results on own-group average participation are not spurious.

In panel E, we present a test based on the opposite idea. In the university, departments are grouped into larger units (such as libraries or the medical school), which we improperly call schools. There are 420 departments distributed among 34 schools in our sample. We regress individual participation on department participation and average participation of other

departments in the same school. A positive coefficient on other departments in the same school would suggest that there is a spurious correlation, at least at the school level. The coefficients of the own department are essentially unaffected (compared to table 3), and the coefficient of the other units within the school are small and insignificant.

The results on participation suggest that the decisions of individuals within one's peer group influence one's decision. They indicate that people may share their knowledge about the plan (and possibly other savings mechanisms). Madrian and Shea (2000a) have replicated these specifications in another context (for a large health insurance company) and find similar results to those we present here. They also present interesting additional evidence. In the firm, all the employees hired after 1998 were automatically enrolled in the TDA. Madrian and Shea (2000b) had previously shown that a very large fraction of these employees remained enrolled in the TDA. Interestingly, they show that the variations in participation across departments caused by the variation in the number of employees enrolled under automatic enrollment *does not* predict larger participation. This is what we expect if workplace effects operate through discussion and information sharing rather than through pure desire for conformity: Presumably the automatic enrollees gave very little thought to the problem, and may even not know that they are enrolled. We would therefore not expect that their enrollment would trigger any enrollment among their peers.

If this interpretation is correct, then we might also expect to see other decisions to be correlated with other people's decisions within the peer group. If employees share information about the plan in the workplace, they presumably also share information about their decisions. For example, asset allocation and the choice of vendor are likely to be influenced by peer decisions. In the next section, we therefore focus on the choice of vendors.

5 Peer Effects on Vendor Choices

We have data on the shares that each participant in the TDA allocates to each vendor. This data is particularly useful, because these vendors offer very similar services at similar prices. Thus, there is less presumption that individuals may be sorted by department according to their intrinsic preference for one vendor versus another. The reader may object that if vendors

are really so similar, then the choice of vendor is of no significance, and therefore there is no particular reason to study it.²¹ However, the very fact that the decision is not of great economic relevance makes it less likely that people are going to feel strongly about it, and therefore that correlated effects are going to be present. In this sense, the results on vendor's choice are interesting *in combination* with the results on participation, because if we find no evidence of peer effects on vendor's choice, it will cast doubt on the results on participation. The reverse is true as well.

Furthermore, even if we were convinced by the results presented in the previous section that there are indeed peer effects in participation, evidence on peer effects in vendor choice will cast some light on the mechanisms through which these effects operate. Peer effects on participation might be driven either by pure peer pressure or by learning effects. Peer pressure arises when individuals imitate the behavior of other people because they think it is the proper thing to do. Other people's behavior informs them of the prevalent norm (as in the case of contraceptive use in Bangladesh, which is studied by Munshi (2000b)). In the case of learning (as in the case of the "green revolution" in India, studied by Foster and Rosenzweig (1995), and Munshi (2000a)), individuals are converging on the optimal economic decision through interactions with their peer group. Peer effects on the decision to save for retirement is probably a result of both norm effects and learning effects. However, peer effects on vendor choice is more likely to be due to learning effects because the choice of a particular vendor probably does not have the status of a prevalent norm. Put another way, if we see that participants in the TDA are influenced by decisions in their workplace about which vendor to choose, it probably means they are thinking that they are learning something about the right economic decisions from conversations with their friends.²²

Table 5 presents the basic results on vendor choice. The dependent variable is the proportion of an individual's TDA contribution which is invested with one of these vendors. The independent regressor of interest is the average of this share for every participant in the department, except the individual.

²¹This issue is, of course, of interest at least for the vendors themselves. More generally, the industrial organization literature is often interested in consumption choices between goods that are close substitutes.

²²This is not to say that they are not all coming to the wrong decision. As Banerjee (1992) shows, this may well happen when people try to infer the correct choice from each other's behavior.

For each vendor, the first column presents the OLS results, the second column presents the 2SLS results, using the proportion of people who fall into each tenure category as instruments.²³ In the third column, we present the 2SLS estimate after controlling for sample selection. We need to control for sample selection because the same factors which make people more likely to participate may make them more likely to choose one vendor rather than another (for example if well informed people both save more and choose a given vendor), which would then introduce a correlation between the unobserved variable in our estimating equation and the equation determining selection in the sample. We control for sample selection using a procedure first suggested by Heckman and Robb (1986), then elaborated by Ahn and Powell (1993), which is to condition on the probability of selection in the sample. Unlike Heckman’s (1979) original procedure of controlling by the Mills ratio, this requires instruments for participation in the TDA which do not influence the choice of vendor conditioning on participation. Fortunately, such instruments are available in this case, since conditional on service, salary influences participation, but not vendor choice. In practice, we first regress individual participation on individual wage, age, gender, and tenure, and the number of individuals who fall in each decile of the university wage distribution in the department (this is the reduced form corresponding to column (2) in table A1). We then calculate the predicted value of participation and the square predicted value, and we include these variables in the first and second stages. Identification is not lost, since we use tenure in the department as an instrument for choice of vendor in the department.

Panel A presents the results in the whole sample, and panel B presents the results from departments where there are at least 15 participants. In panel A, for the three vendors, the OLS coefficients are smaller than the IV coefficients. They are very small and statistically not different from 0 in the case of vendor D and vendor V. In contrast, the IV coefficient are large and significant. They are also very similar from vendor to vendor. They indicate that when the share invested in a given vendor by other contributors in one’s department raises by 1 percentage point, one’s share of investment with this vendor raises by half a percentage point. Controlling for sample selection does not affect this coefficient. This suggests that the

²³Salary is not a valid instrument in this case because it does not predict vendor choice. The F-statistic of the first stage for the choice of vendor are large, especially for vendors R and V. Vendor V was introduced later, and there is a strong negative correlation between tenure and the share allocated to vendor V. In contrast, there is a strong positive correlation between year of service and the share allocated to vendor R.

propensity to save is not systematically correlated with vendor choice. This confirms that the factors which influence participation do not at the same time influence vendor choice. Therefore, it gives us some confidence that individuals' preferences for a particular vendor are unlikely to differ systematically from one department to another, even if propensity to participate does.

In Panel B, we restrict the sample to departments where at least 15 staff employees participate in the TDA. The OLS coefficients in Panel B are much larger than in Panel A whereas the IV coefficients of Panel B are slightly smaller than in Panel A. This pattern is similar to what we found for participation.

Table 6 presents the 2SLS results on the effects of participation of department members broken down by groups. We cannot present any results broken down by tenure, since tenure is our only instrument. However, we present results by age, gender, and faculty. Again, cross-group effects are absent in all cases and for all vendors, while own-group effects are positive and large in all cases, except for the men in the case of vendor R. The patterns are therefore comparable to what we have seen in the case of participation.

6 Conclusion

In this paper, we set out to study the role of peer group effects on the decision to participate in the TDA and on the choice of vendor among participants. Identifying endogenous social effects is almost an impossible task in most cases where assignment to a peer group is not random. Most individuals' decisions within a social group are correlated for reasons which have nothing to do with the fact that individuals are imitating each other. Their decisions may be influenced by common variables, observed or unobserved, such as taste, background, or common environmental factors. The application studied in this paper is a favorable case, since individuals in the university share the same plan and the same program inputs. An important source of correlation between individual's behavior is therefore eliminated.

We recognize, however, that the participation of individuals within departments may be correlated because they may share a common propensity to save. After instrumenting average participation in the department with the distribution of wages in the department or the distribution of years of service, a strong effect of average participation within sub-groups in a

department (along gender, service, status, or age lines) persists. In contrast, we find no effect of the participation in the other sub-group within the department. The same results are obtained for the choice of vendor among participants to the TDA. We interpret these results as very suggestive evidence that decisions taken in one's peer group influence one's decision to participate and the choice of the mutual fund vendor. When participation increases by 1 percent in the department, one's participation increases by 0.15 percent. When the average share of the contribution invested in one vendor increases by 1 percent, one's share in this vendor increases by 0.5 percent on average.

These results, if confirmed, have several important implications. First, they contribute to the literature on the determinants of retirement savings. The work of Bernheim on financial education ((Bernheim and Garrett 1996)), and Madrian and Shea (2000b) on default rules have shown that economic incentives are not the only determinants of savings decisions. This paper adds to these studies by showing that peer effects are another source of extra-economic influence on people's decisions. Individuals do not instantly learn about economic opportunities, and their environment is a strong determinant of their economic decisions. Low levels of savings by American households have been a source of preoccupation for academics and policy makers alike. Recognizing that savings decisions are influenced by peer's savings decisions could be an important element to improve our understanding of these issues. More generally, recognizing that the financial decisions of a majority of people are influenced by the actions of others should be an important element in the way we incorporate individual decisions into macroeconomic models. In future research, we plan to study the role of peer effects in the allocation of assets.

Second, these results provide a possible rationale for organizing 401(k)s around the workplace. In the case of tax deferred accounts which individuals can access on their own and outside the workplace (such as IRAs), people have no obvious peer group with which to discuss their choices. The strong decline in participation in IRAs following the Tax Reform Act of 1986 has been considered as evidence that advertisement and information are one of the key elements driving participation rates (see Bernheim (1999)). When the TDA is organized by employers such as in the case of 401(k) plans, co-workers become a natural group with which to discuss it as the benefits package is common to employees, and therefore a likely conversation topic. Offering savings options organized around the workplace may therefore increase the overall level

of savings.²⁴

In this paper, we make no attempt to distinguish the effect of behavior from the effect of characteristics which led to this behavior. Assuming that our results can be interpreted as evidence that the savings behavior of my colleagues affects my saving behavior, an important question remains unanswered. Is it because I am influenced by their *characteristics* or by their *actions* (in Manski's words, is the social effect a 'contextual' effect or an 'endogenous' effect)? In other words, do I save more because my colleagues are the type of people who save, or because I see them saving and I imitate them? This distinction has strong policy implications because it determines whether or not there could be a 'multiplier effect' of financial education. If endogenous effects are important, the role of financial education may go far beyond providing information to those who are directly exposed to it. If a few individuals enroll in the plan following an information session, it might trigger non-negligible repercussion effects. This effect is potentially important when assessing the effect of education or information sessions on contribution decisions in voluntary retirement plans such as 401(k)s.

We, therefore, see this study as a first step in a broader research agenda. A simple look at existing, non-experimental data has suggested that peer effects may be present and important. This was confirmed by a few informal conversations with employees. In future work, we plan to address the shortcomings of the present study and to disentangle contextual and endogenous social effects. We plan to explore two directions. First, we would like to administer an in-depth survey to a sample of employees in the university, in which we would ask about the key sources of information that have induced individuals to enroll (or not) in the plan. Second, we are planning a randomized financial education experiment, which might alter the participation rate of a random subset of employees within a subset of randomly chosen departments. We will then compare the participation rate of the non-affected individuals in these departments with the participation rate in the departments where no one was affected. By doing so, we will directly address the question of the multiplier of financial education efforts.

²⁴Although in principle it could go either way.

References

- Banerjee, Abhijit V. (1992) 'A simple model of herd behavior.' *Quarterly Journal of Economics*
- Bayer, Patrick, B. Douglas Bernheim, and Karl Scholz (1998) 'The effects of financial education in the workplace: Evidence from a survey of employers.' Working Paper, National Bureau of Economic Research
- Bernheim, B. Douglas (1994) 'A theory of conformity.' *Journal of Political Economy* 102(5), 841–877
- (1999) 'Taxation and saving.' Working Paper 7061, National Bureau of Economic Research, March
- Bernheim, B. Douglas, and Daniel M. Garrett (1996) 'The determinants and consequences of financial education in the workplace: Evidence from a survey of households.' Working Paper 5667, National Bureau of Economic Research, July
- Bertrand, Marianne, Sendhil Mullainathan, and Erzo Luttmer (1998) 'Network effects and welfare cultures.' Working Paper 6832, National Bureau of Economic Research, December
- Besley, Timothy, and Anne Case (1994) 'Diffusion as a learning process: Evidence from HYV cotton.' Discussion Paper 174, RPDS, Princeton University
- Case, Anne, and Larry Katz (1991) 'The company you keep: The effect of family and neighborhood on disadvantaged youths.' Working Paper 3708, National Bureau of Economic Research
- Elison, Glenn, and Drew Fudenberg (1993) 'Rules of thumbs for social learning.' *Journal of Political Economy* 101(4), 93–126
- Engen, Eric M., William G. Gale, and John K. Scholz (1996) 'The illusory effect of saving incentives on saving.' *Journal of Political Economy* 10(4), 113–138
- Evans, William, Willian Oates, and Robert Schwab (1992) 'Measuring peer group effects: A model of teenage behavior.' *Journal of Political Economy* 100(5), 966–991

- Foster, Andrew D., and Mark R. Rosenzweig (1995) 'Learning by doing and learning from others: Human capital and technical change in agriculture.' *Journal of Political Economy* 103(6), 1176–1209
- Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman (1996) 'Crime and social interactions.' *Quarterly Journal of Economics* CXIV(2), 502–548
- Kusko, Andrea, James M. Poterba, and David W. Wilcox (1994) 'Employee decisions with respect to 401(k) plans: Evidence from individual level data.' Working Paper 4635, National Bureau of Economic Research
- Madrian, Brigitte, and Dennis Shea (2000a) 'Peer effects and savings behavior in employer-sponsored savings plans.' Mimeo, University of Chicago
- (2000b) 'The power of suggestion: Inertia in 401(k) participation and savings behavior.' Mimeo, University of Chicago
- Manski, Charles (1993) 'Identification of exogenous social effects: The reflection problem.' *Review of Economic Studies* 60, 419–431
- Munshi, Kaivan (2000a) 'Social norms and individual decisions during a period of change : An application to the demographic transition.' MIMEO
- (2000b) 'Social norms, with an application to the demographic transition in Bangladesh.' MIMEO, University of Pennsylvania
- Papke, Leslie E. (1994) 'Participation in and contributions to 401(k) plans: Evidence from plan data.' *Journal of Human Resources*
- Papke, Leslie E., Mitchell Petersen, and James M. Poterba (1993) 'Did 401(k) plans replace other employer provided pensions?' Working Paper 4501, National Bureau of Economic Research, October
- Poterba, James M., Steven F. Venti, and David A. Wise (1996) 'How retirement saving programs increase saving.' *Journal of Economic Perspectives* 10(4), 91–112

Sacerdote, Bruce (2000) 'Peer effects with random assignment: Results for Dartmouth roommates.' Working Paper 7469, National Bureau of Economic Research

Table 1: Participation Rates in the University Libraries

Rank by Participation Level	Participation Rate	Average Wage	Average Tenure	Number of Staff
(1)	(2)	(3)	(4)	(5)
1	0.73	\$33,600	12.9	22
2	0.50	\$36,300	11.3	28
3	0.47	\$34,500	11.2	15
4	0.44	\$33,100	9.4	16
5	0.43	\$34,600	12.7	7
6	0.38	\$29,400	8.7	13
7	0.38	\$39,100	11.6	29
8	0.34	\$36,200	11.0	235
9	0.32	\$32,100	8.8	47
10	0.23	\$34,600	10.2	17
11	0.14	\$31,600	9.8	7

Notes: Column (2) displays the participation rate in the 403(b) plan in the 11 libraries of the university library system. Columns (3) and (4) report average wages and tenure, respectively, in each library. Column (5) reports the number of staff employees in each library.

Table 2: Summary Statistics

	Complete Sample	Staff	Faculty
	(1)	(2)	(3)
Panel A: General Characteristics			
Number of employees, 10/97	10,661	8,266	2,362
Number of employees, 06/98	10,762	8,443	2,303
Number of employees, 10/98	10,876	8,627	2,214
Number of employees, 06/99	11,351	9,150	2,181
Salary (mean)	\$48,330 [31,930]	\$39,200 [22,030]	\$83,020 [42,140]
Salary (median)	\$38,300	\$34,500	\$72,100
Gender (Male)	0.45 [0.50]	0.39 [0.49]	0.71 [0.45]
Age	41.8 [11.7]	40.0 [11.1]	48.9 [11.6]
Tenure (years of service)	8.1 [8.8]	7.3 [7.8]	11.6 [11.3]
Panel B: TDA Participation and Vendor Choices			
Participation in the TDA plan	0.352 [0.477]	0.327 [0.476]	0.448 [0.497]
Participation in the TDA plan when Tenure<3 months	0.034 [0.181]	0.030 [0.169]	0.065 [0.239]
Share of Vendor R	0.351 [0.455]	0.342 [0.451]	0.379 [0.466]
Share of Vendor D	0.325 [0.429]	0.339 [0.434]	0.286 [0.412]
Share of Vendor V	0.221 [0.383]	0.220 [0.378]	0.223 [0.395]
Panel C: Departments			
Number of Departments	420		
Average Size of Departments	32 [51]		
Median Size of Departments	12		

Notes: Food and custodial services excluded. Business school also excluded.

Table 3: Peer Effects on Participation Decisions for Sub-samples

Instruments	OLS	2SLS Salary deciles	2SLS Tenure dummies	2SLS Tenure Salary
	(1)	(2)	(3)	(4)
PANEL A. STAFF				
Average participation in the department	0.271 (0.035)	0.130 (0.079)	0.134 (0.081)	0.149 (0.062)
Number of Observations	29,863	29,863	29,863	29,863
PANEL B: DEPARTMENTS WITH LESS THAN 20 EMPLOYEES				
Average participation in the department	0.210 (0.052)	0.234 (0.137)	0.347 (0.136)	0.300 (0.109)
Number of Observations	4,520	4,520	4,520	4,520
PANEL C: DEPARTMENTS WITH MORE THAN 20 EMPLOYEES				
Average participation in the department	0.317 (0.047)	0.087 (0.090)	0.137 (0.093)	0.114 (0.070)
Number of Observations	25,377	25,377	25,377	25,377

Notes: Standard errors are corrected for clustering.

The regressions include all the control variables listed in table A1.

In column 2, the instruments are the proportion of employees of the department whose wage falls into each decile of the university-wide distribution of wages. In column 3, the instruments are the number of employees of the department whose tenure falls into each category.

Table 4: Peer Effects on Participation Decisions Among Sub-groups

Instruments	Group 1		Group 2	
	OLS	2SLS	OLS	2SLS
	(1)	Salary Tenure (2)	(3)	Salary Tenure (4)
PANEL A: GROUP 1: FEMALE (17,493 obs.) AND GROUP 2: MALE (11,066 obs.)				
Average participation in group 1 (female)	0.199 (0.045)	0.321 (0.085)	0.093 (0.052)	-0.053 (0.092)
Average participation in group 2 (male)	0.068 (0.040)	-0.115 (0.071)	0.121 (0.049)	0.275 (0.085)
p-value of test coeff differ	0.055	0.001	0.727	0.029
PANEL B: GROUP 1: TENURE BELOW 3 YEARS (10,350 obs.) AND GROUP 2: ABOVE 7 YEARS (12,307 obs.)				
Average participation in group 1 (young)	0.114 (0.038)	0.235 (0.083)	-0.008 (0.048)	-0.107 (0.106)
Average participation in group 2 (old)	0.063 (0.034)	0.048 (0.077)	0.188 (0.053)	0.274 (0.139)
p-value of test coeff differ	0.345	0.108	0.008	0.034
PANEL C: GROUP 1: YOUNG (35 AND BELOW, 11,358 obs) AND GROUP 2: OLD (ABOVE 36, 17,269 obs)				
Average participation in group 1 (young)	0.218 (0.042)	0.396 (0.075)	0.031 (0.042)	-0.059 (0.073)
Average participation in group 2 (old)	0.023 (0.044)	-0.206 (0.088)	0.244 (0.047)	0.287 (0.092)
p-value of test coeff differ	0.003	0.000	0.003	0.011
PANEL D: GROUP 1: STAFF (23,304 obs.) AND GROUP 2: FACULTY				
Average participation in group 1 (female)	0.187 (0.039)	0.128 (0.057)		
Average participation in group 2 (male)	0.016 (0.024)	-0.073 (0.045)		
p-value of test coeff differ	0.000	0.008		
PANEL E: GROUP 1: DEPARTMENT AND GROUP 2: OTHER DEPARTMENTS IN THE SAME SCHOOL				
Average participation in group 1 (own department)	0.262 (0.036)	0.164 (0.063)		
Average participation in group 2 (other departments)	0.027 (0.042)	0.023 (0.042)		
p-value of test coeff differ	0.000	0.061		

Notes: Standard errors are corrected for clustering. The regressions include all the control variables listed in table A1. Instruments are the proportion of employees of the department whose wage falls into each decile of the university-wide distribution of wages and the number of employees of the department whose tenure falls into each category. In Panel B, only the salary instruments are used.

Table 5: Peer Effects on Vendor Choice

	Share of vendor R			Share of vendor D			Share of vendor V		
	OLS	2SLS	2SLS	OLS	2SLS	2SLS	OLS	2SLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: All Departments with Number of Participants above 2									
Average share of vendor among participants in the department	0.262 (0.047)	0.511 (0.132)	0.542 (0.137)	0.073 (0.052)	0.633 (0.275)	0.73 (0.292)	-0.084 (0.042)	0.481 (0.218)	0.467 (0.209)
F-statistic of the first stage		10.3	9.8		3.4	3.0		9.6	8.4
Number of Observations	10,233	10,233	10,233	10,233	10,233	10,233	10,233	10,233	10,233
Control for sample selection	No	No	Yes	No	No	Yes	No	No	Yes
Panel B: Departments with Number of Participants above 15									
Average share of vendor among participants in the department	0.507 (0.072)	0.511 (0.146)	0.539 (0.140)	0.313 (0.098)	0.44 (0.281)	0.478 (0.271)	0.071 (0.093)	0.468 (0.271)	0.489 (0.262)
F-statistic of the first stage		6.8	8.1		2.9	3.6		6.7	5.9
Number of Observations	7,584	7,584	7,584	7,584	7,584	7,584	7,584	7,584	7,584
Control for sample selection	No	No	Yes	No	No	Yes	No	No	Yes

Notes: Standard errors (corrected for clustering at the individual levels) are in parentheses. All regressions control in addition for gender, salary deciles, and age dummies. The excluded instrument in the equation predicting the participation are the proportion of members of the departments who fall in each category of years of service.

Table 6: Peer Effects on Vendor Choice Among Sub-groups

Instruments	Young (below 45) and Old (above 45)		Male and Female		Faculty and Staff
	Group 1 (Young)	Group 2 (Old)	Group 1 (Female)	Group 2 (Male)	Group 1 (Staff)
	2SLS Service dummies	2SLS Service dummies	2SLS Service dummies	2SLS Service dummies	2SLS Service dummies
	(1)	(2)	(3)	(4)	(5)
PANEL A1: Vendor R					
Average participation in group 1	0.87 (0.358)	0.037 (0.444)	0.681 (0.195)	0.264 (0.225)	0.44 (0.186)
Average participation in group 2	0.025 (0.223)	0.605 (0.486)	-0.141 (0.241)	0.114 (0.381)	0.098 (0.106)
PANEL A2: Vendor D					
Average participation in group 1	1.483 (0.522)	0.426 (0.625)	0.812 (0.265)	0.124 (0.267)	0.833 (0.195)
Average participation in group 2	-0.284 (0.259)	0.707 (0.256)	-0.26 (0.193)	0.331 (0.308)	-0.094 (0.112)
PANEL A3: Vendor V					
Average participation in group 1	0.607 (0.265)	0.021 (0.256)	0.597 (0.229)	-0.162 (0.267)	0.593 (0.238)
Average participation in group 2	0.076 (0.284)	0.683 (0.438)	-0.177 (0.240)	0.393 (0.434)	-0.083 (0.122)
Number of Observations	4,555	4,576	5,974	3,172	7,252

Notes: Standard errors (corrected for clustering at the individual levels) are in parentheses. All regressions control in addition for gender, salary deciles, and age dummies. The excluded instrument in the equation predicting the participation are the proportion of members of the departments who fall in each category of years of service.

Table A1: Coefficients for All Variables and Controls

	OLS	2SLS	2SLS	2SLS
		Salary	Tenure	Tenure and Salary
	(1)	(2)	(3)	(4)
Average participation in the department	0.271 (0.035)	0.130 (0.079)	0.134 (0.081)	0.149 (0.062)
Gender (male dummy)	-0.063 (0.010)	-0.067 (0.010)	-0.066 (0.010)	-0.066 (0.010)
Age				
10 to 19	0.137 (0.030)	0.134 (0.029)	0.134 (0.029)	0.134 (0.029)
20 to 29	0.193 (0.023)	0.189 (0.023)	0.190 (0.023)	0.189 (0.023)
30 to 39	0.190 (0.021)	0.185 (0.022)	0.186 (0.022)	0.186 (0.022)
40 to 49	0.200 (0.022)	0.196 (0.022)	0.196 (0.022)	0.197 (0.022)
50 to 59	0.282 (0.023)	0.279 (0.023)	0.279 (0.023)	0.279 (0.023)
60 to 69	0.347 (0.032)	0.342 (0.032)	0.343 (0.032)	0.343 (0.032)
70 to 79	0.142 (0.059)	0.139 (0.059)	0.139 (0.059)	0.139 (0.059)
Tenure				
Less than 1 years	-0.174 (0.019)	-0.178 (0.019)	-0.178 (0.019)	-0.177 (0.018)
1 to 2 years	-0.096 (0.020)	-0.097 (0.020)	-0.097 (0.020)	-0.097 (0.020)
3 to 4 years	-0.039 (0.020)	-0.038 (0.020)	-0.038 (0.020)	-0.038 (0.020)
4 to 7 years	-0.031 (0.020)	-0.030 (0.021)	-0.030 (0.020)	-0.030 (0.020)
7 to 12 years	-0.002 (0.019)	-0.002 (0.019)	-0.001 (0.019)	-0.002 (0.020)
Salary decile				
Decile 1	-0.248 (0.029)	-0.252 (0.029)	-0.252 (0.029)	-0.252 (0.029)
Decile 2	-0.226 (0.029)	-0.231 (0.029)	-0.230 (0.029)	-0.231 (0.029)
Decile 3	-0.202 (0.028)	-0.205 (0.028)	-0.205 (0.028)	-0.205 (0.028)
Decile 4	-0.149 (0.029)	-0.152 (0.029)	-0.152 (0.030)	-0.151 (0.030)
Decile 5	-0.085 (0.030)	-0.087 (0.030)	-0.087 (0.030)	-0.086 (0.030)
Decile 6	-0.004 (0.029)	-0.007 (0.029)	-0.007 (0.029)	-0.007 (0.029)
Number of Obs.	29,863	29,863	29,863	29,863
F-Statistic of the First Stage		15.8	28.9	26.7

Notes: Sample restricted to Staff employees. Food service and business school employees excluded.

Standard errors corrected for clustering at the individual levels.

In column 2, the instruments are the proportion of employees of the department whose wage falls into each decile of the university-wide distribution of wages. In column 3, the instruments are the number of employees of the department whose tenure falls into each category. In column 4, both sets are used together.