How to Make the Team: Social Networks vs. Demography as Criteria for Designing Effective Teams

Ray Reagans
Columbia University
Columbia Business School

Ezra Zuckerman
Massachusetts Institute of Technology
Sloan School of Management

Bill McEvily
Carnegie Mellon University
Graduate School of Industrial Administration

Abstract

We compare two alternative approaches for evaluating the potential of a work group or team: one that focuses on team members' demographic characteristics and one that focuses on the members' social networks. Given that people's network contacts often share their demographic attributes (i.e., the network is homophilous), the two approaches seem equivalent, and the first seems preferable because it is easier to implement. In this paper, we demonstrate several important limits to this rationale. First, we argue and show, in an analysis of 1,518 project teams in a contract research and development firm, that even when internal organizational networks are significantly homophilous with respect to demographic variables, the very logic of the causal structure that underlies theories of demographic diversity carries ambiguous performance implications. This ambiguity is due to the fact that demographic diversity has opposing effects on two social network variables—internal density and external range—that each has a positive effect on a teams' performance. We also demonstrate that a focus on demographic criteria is problematic because the demographic makeup of an organization can place inherent limits on a manager's ability to shape the demographic composition of a team. The ambiguous performance implications and the inherent limits placed on a manager's ability to manage a teams' demography reduce the likelihood that a manger's interventions will be successful. The performance implications of managing a team's social capital, however, are clear.

The chief challenge for a manager in assigning people to a project team is to evaluate their potential for helping the team achieve high performance. If a key factor in determining team performance is its capacity to access a wide range of the information, resources, or perspectives that are distributed throughout an organization, a manager or team leader would need to maximize such breadth of access. One approach to doing so is to assess prospective team members based on their demographic characteristics. For example, if employees who enter the organization at proximate points in time have similar experiences and perspectives, one might expect higher performance from teams whose members are drawn from various cohorts because such teams have access to a wider array of knowledge and information, which should increase their capacity for learning and creative problem solving. An alternative approach for maximizing breadth of access is to focus on social networks rather than demographic characteristics. Because information, resources, and perspectives tend to be more similar within than between network clusters, selecting team members who belong to various clusters in the organization's social network would provide the team with a greater potential for success.

The organizational demography literature (see Williams and O'Reilly, 1998, for a review) has tended to assume that these alternate approaches are interchangeable (Lawrence, 1997). The rationale for doing so is based on the assumption that the social network inside the organization (or "organizational network") is characterized by homophily, the tendency for strong network connections to occur more frequently between people who share an important demographic characteristic such as race, gender or firm tenure (see McPherson, Smith-Lovin, and Cook, 2001 for a review; cf. Laumann, 1966; Byrne, 1971; Blau, 1977; Blau and Schwartz, 1984; Brass, 1985; Zenger and Lawrence, 1989; Ely, 1994). Insofar as the tendency toward homophily is strong in a particular organization, demographic criteria and network-based criteria are equivalent. For example, if ties are much stronger within than between organizational cohorts, staffing a team based on its tenure distribution produces the same results as selecting team members based on their networks, and vice versa. Moreover, if demographic and network-based criteria are interchangeable, there is good reason to prefer the former, since demographic data are less difficult for both the researcher and the manager to collect. As Pfeffer (1982, 1983) argued, a focus on demography makes sense even if we assume that the effect of demographic composition on performance is not causal but either reflects or is mediated by intervening processes such as network-based dynamics. Since demographic variables are easier to "access and reliably measure," and if networks are strongly influenced by demography, managers and researchers may profitably concentrate on the former (Pfeffer, 1982: 351; cf. Tsui and Gutek, 1999). Insofar as the demographic composition of a team is a useful if imperfect proxy for underlying social-network-based factors that determine project performance, focusing on team demography seems appropriate, especially since this practice appears inexpensive to use when compared with collecting detailed social network data.

There are, however, strong reasons to question this justification for the use of demographic rather than social network-based criteria for anticipating team performance. Lawrence (1997) has cast doubt on the assumption that organizational networks are characterized by a significant amount of homophily. She reviewed evidence indicating that observed levels of homophily are generally too weak to justify using demographic variables as proxies for network variables. The available empirical evidence indicates that a variety of contextual factors can reduce the degree of homophily inside an organization or team. In particular, if the team members have past experience interacting with members of different social categories (Westphal and Milton, 2000) or if the organization culture fosters the creation of cross-category relationships (Chatman et al., 1998; Ely and Thomas, 2001; Polzer, Milton, and Swann, 2002), then homophily will be attenuated. In general, sharing a demographic characteristic produces a strong network connection the more that people identify with their shared characteristic (Mehra, Kilduff, and Brass, 1998). As the level of identification declines, so does the tendency for strong ties to be concentrated among those who possess the attribute.

Yet the recognition that homophily may be weak in certain organizations is not sufficient to rule out the use of demography as a guide for staffing teams. Some organizations are marked by a significant amount of homophily, at least with respect to particular demographic attributes, and in these organizations, emphasizing demographic diversity would seem to be useful. Moreover, even if homophily is weak, demographic variables might still recommend themselves as relatively inexpensive, if imperfect, proxies for social network variables. Yet there are several key difficulties with relying on a demography-based approach beyond the reason raised by Lawrence (1997). Most notably, <u>even when social networks are homophilous with respect to a given demographic attribute</u>, demographic diversity has dual, contradictory implications for the network-based mechanisms that shape team performance. Reagans and Zuckerman (2001) showed that there is general consensus about the overall causal model linking demographic diversity to network-based social capital variables and linking those network variables to team performance. Yet despite consensus on the causal structure, which is reproduced in figure 1, it remains unclear how diversity affects performance. As depicted in the figure, demographic diversity is hypothesized to <u>decrease</u> a team's internal density (team members of different demographic categories are presumed to have relatively weak relationships with one another) and to increase the amount of range in the team's external network (team members of different demographic categories are presumed to be able to reach different constituencies outside the team), and each network variable is expected to have a positive effect on team performance. Since demographic diversity has opposing effects on two key network variables and those network variables both have positive effect on team performance, it is not surprising that the available empirical evidence indicates that greater demographic diversity may be associated with increases, decreases, or no corresponding change in team performance (Williams and O'Reilly, 1998). Given this causal structure, it would be more surprising if there were a consistent association between demographic diversity and team performance.

FIGURE 1 ABOUT HERE

These considerations carry a key empirical implication, the validation of which forms the primary objective of this paper: that it is possible for both social network variable to have significant effects on team performance but for changes in demographic diversity to have no net effect because diversity has opposing effects on each of these network variables. Our empirical analysis draws on data from a contract research and development (R&D) firm that are especially well-suited for validating the full causal structure depicted in figure 1 because they allow for direct measurement of social networks both inside and outside a team (cf., Reagans and Zuckerman 2001). These data also allow us to highlight a second difficulty with the demography-based approach. In particular, in order to staff a team effectively, a manager must recognize the factors that limit his ability to manipulate a team's demographic composition. There are two basic reasons for such limits. First, the nature of the project typically necessitates the inclusion of certain types of employees either for functional or political reasons. Second, the various demographic characteristics are often correlated with one another, thereby restricting a manager's ability to manipulate one demographic characteristic while holding others constant. Accordingly, before modeling the two stages of figure 1, we first analyze the demographic landscape of the contract R&D organization we study and show how that landscape inherently curtails the degree to which it is possible for a manager to shape the demographic composition of a team. Before introducing the data and proceeding with such analyses, we review the theoretical basis for the causal pathways depicted in figure 1.

TEAM STAFFING AND PERFORMANCE

Causal Structure

Reagans and Zuckerman (2001) provided an extensive discussion of the causal linkages between demographic diversity and team performance, as portrayed in figure 1. Their key insight is that each of the two opposing perspectives in the debate about the performance implications of demographic diversity relies on one of the two network-based mechanisms discussed in the social capital literature: network density and network range. Further, since these two mechanisms may operate independently, the apparent opposition between them is largely illusory. Those who are pessimistic about the performance implications of diversity (e.g., McCain, O'Reilly, and Pfeffer, 1985; O'Reilly, Caldwell, and Barnett, 1989; Zenger and Lawrence, 1989) tend to stress the importance of network density for reasons that are similar to the arguments advanced by scholars who stress the importance of social network "closure," or tightly connected networks for providing social capital. A common metaphor in both literatures is the mutual support and coordination often found in cohesive, ethnic communities (e.g., Aldrich and Zimmer, 1986; Portes and Sensenbrenner, 1993). Network closure in such communities is thought to facilitate mutual identification among members of a collectivity (e.g., Portes and Sensenbrenner, 1993) and to promote a degree of trust sufficient to support social exchange and collective action (Coleman, 1988). To the extent that closure promotes the development and enforcement of norms and insofar as these norms emphasize high performance, increasing network closure can be expected to have a positive effect on performance (e.g., Homans, 1950). Thus, assuming such norms are present, pessimists see organizational diversity as problematic because diverse teams are unlikely to assume a cohesive, community-like character (Pfeffer, 1985; McCain, O'Reilly, and Pfeffer, 1985; O'Reilly, Caldwell, and Barnett, 1989). In this view, demographic diversity reduces internal coordination, which hinders a team's ability to succeed ($\lambda_2 > 0$ in figure 1).

While the absence of network closure -- or the presence of "structural holes" between contacts-- is problematic within a team (Cummings and Cross, 2003; Oh, Chung, and Labianca, 2003), a second strand of the social capital literature focuses on the advantages of structural holes between contacts. The social capital provided by closure is thought to come at the expense of the social capital provided by structural holes. Yet Reagans and Zuckerman (2001) clarified that each applies at different locations in the social structure (see also Burt, 2000; Gabbay and Zuckerman, 1998; Ingram and Roberts, 2000; Reagans and McEvily, 2003; cf. Burt 1980; 1982, ch. 7; 1992, ch. 3). Figure 2 distinguishes between internal (or local) and external (or global) structural holes. While local structural holes, like that between members A and B, hinder internal coordination and the team's capacity for collective action, ties that bridge across holes outside the team (like that between the contacts networks of members A and C) generate "information benefits" because they represent points of contact into different network clusters, each of which tends to represent a relatively non-redundant concentration of information and resources (Burt 1992: 13-16; Reagans, Zuckerman, and McEvily, 2004). Such boundary-spanning ties provide access to a broader array of ideas and opportunities than do ties that are restricted to a single cluster (Granovetter, 1973). Researchers who are optimistic about diversity's effect on performance ground their views in this logic. According to optimists, teams that draw members from diverse demographic categories benefit because such teams generate links between people with different information, resources, and perspectives. The greater range associated with diversity enhances the team's capacity for learning and creative problem solving (Bantel and Jackson, 1989; March, 1991; Ancona and Caldwell 1992; Pelled, Eisenhardt, and Xin, 1999; Reagans and McEvily, 2003). By introducing divergent thoughts and opinions, external range enhances performance ($\lambda_1 > 0$ in figure 1).

[FIGURE 2 ABOUT HERE]

Thus, the two network effects we have reviewed ($\lambda_1 > 0$ and $\lambda_2 > 0$ in figure 1) do not contradict one another, since the social capital arguments that support them apply to different locations in social structure and address two distinct challenges identified by research on teams (see Hackman and Morris, 1975): creative problem solving (facilitated by external networks that bridge structural holes) and implementation and execution (facilitated by strong internal networks). And participants in debates about the effects of demographic diversity tend to argue not that one effect is valid and the other is invalid but that, depending on one's focus, one outweighs the other (e.g., Pfeffer, 1985; Ancona and Caldwell, 1992).

Moreover, there is similar agreement on the valence and significance of the causal paths that link demographic diversity with the two network variables, internal density and external range. In particular, diversity is thought to reduce internal density ($\beta_2 < 0$ in figure 1) while it increases external range ($\beta_1 > 0$ in figure 1). The rationale for these predictions is grounded in social psychological and sociological theories of homophily (McPherson, Smith-Lovin, and Cook, 2001). Social psychological explanations for homophily rely on self-categorization and social identity theory (Tajfel, 1981; Tafjel and Turner, 1986; Turner, 1987). Increasing demographic diversity makes an individual more aware of his or her social characteristics, and the individual identifies more with people who share his or her characteristics while ascribing negative characteristics and motives to out-group members (Brewer, 1979; Kramer, 1991). Sociological approaches to homophily sometimes emphasize similar preferences (e.g., Laumann, 1966; Byrne, 1971) but also identify environmental and structural conditions that produce differential tendencies for individuals from different backgrounds to come into contact and therefore to develop strong ties (Blau 1977; Feld, 1982; Blau and Schwartz, 1984).

While scholars differ on the factors they cite as producing homophily, the implications of homophily for the organization of informal networks and thus for the internal and external networks of teams are clear. At the extreme, people who share an attribute will be strongly connected to each other and to the same people (White, 2002). And structural holes should be observed between people who do not share an attribute. In such a segregated network, increasing demographic diversity inside a team promotes range in the team's external network because the team is composed of individuals who are connected to different external contacts. As demographic diversity increases on a team, so does the amount of range in the team's external network ($\beta_1 > 0$ in figure 1). Conversely, a homogenous team is more likely to be embedded in a dense web of third-party interconnections because socially similar individuals are more likely to share the same contacts. Therefore, increasing demographic diversity restricts internal density ($\beta_2 < 0$ in figure 1).

Although there is consensus on the causal structure that links demographic diversity with internal density and external range and, indirectly, with performance, there seems to be less appreciation for what this causal structure implies for debates about the management of diversity. In particular, if we assume that each pair of effects at each point in the causal chain is of the same magnitude (i.e., $|\lambda_1| = |\lambda_2|$ and $|\beta_1| = |\beta_2|$), then the overall effect of demographic diversity on performance will be nil even if the organizational network is marked by significant homophily (i.e., $\beta_1 >> 0$ and $\beta_2 << 0$). In such circumstances, each pathway of the causal chain will equal the reverse of the other ($\beta_1\lambda_1 = -\beta_2\lambda_2$), thereby canceling each other out (i.e., $\delta = \beta_1\lambda - \beta_2\lambda_2 = 0$). Thus to justify an exclusive focus on team demography even when homophily is strong, one must relax the assumption that each pair of effects is of equal magnitude. Yet insofar as the extent of homophily governs <u>both</u> the effect of diversity on internal density and its effect on external range, there is little reason to believe that the first pair of effects will not cancel each other out. In addition, it is not clear why, as a general tendency, the effect of external range on performance might be greater (or less) than that for internal density, as optimists (or pessimists) appear to assume.

Without a strong basis for privileging one causal pathway over the other, demographic diversity constitutes an ambiguous criterion for judging a project team's potential for high performance <u>even under conditions of significant homophily with respect to the demographic attribute</u>. That is, not only is research on demography and performance problematic when the assumption of homophily cannot be supported (Lawrence, 1997), it is also problematic under conditions of significant homophily. Thus, our empirical analysis below is designed to validate the entire causal structure while leaving open the possibility that there is no net association between demographic diversity and performance.

Limits to Managerial Discretion

Before proceeding with these analyses, we consider how the limits to managerial discretion restrict the applicability of the demography-based approach. To staff a team effectively, a manager must recognize the factors that limit his or her ability to manipulate a team's demographic composition. There are two basic reasons for such limits. First, the nature of the project typically necessitates the inclusion of certain types of employees either for functional or political reasons. Such requirements thus restrict the scope for choice regarding demographic composition. Second, the various demographic characteristics are often correlated with one another, thereby restricting a manager's ability to manipulate one demographic characteristic, while holding the other demographic characteristics constant. In practice, the assumption that demographic variables are orthogonal to one another can produce unintended outcomes. For example, if an organization's internal network is characterized by considerable gender homophily, and the organization only recently began hiring women, maximizing a team's tenure-based diversity will tend to lower the team's gender diversity. To clarify this issue, let us assume that the degree of homophily and the relevant causal relationships in figure 1 come to be understood somehow, perhaps through trial-and-error learning. To what extent can the manager then manipulate a team's demographic composition in an attempt to maximize its performance?

It is useful to distinguish first between two different kinds of characteristics: those that are fixed, in the sense that they reflect the task or functional requirements of the project, and those that are variable, in the sense that they relate to resources that can be made available to the project in varying proportions. Put differently, in allocating or recruiting personnel to a project, there are certain attributes that must be represented (e.g., expertise in mechanical engineering) and thus limit the manager's discretion in designing the team. Within these fixed parameters, the manager typically enjoys some degrees of freedom in selecting from among the available personnel that meet those parameters (e.g., individuals of either high or low tenure).

The distinction between fixed and variable characteristics becomes even more salient when we consider a second issue that complicates the demography-based approach: managers must consider multiple demographic characteristics simultaneously, and these attributes vary in the extent to which they jointly produce multiform homogeneity or heterogeneity (Blau, 1977: 77). Multiform homogeneity refers to the extent to which demographic characteristics are consolidated such that membership in one implies membership in the other. For example, race and gender are orthogonal in society because members of each race are equally likely to be male and female. By contrast, race and religion are somewhat consolidated in the U.S. in that, for instance, most African Americans are Protestants (and members of particular denominations). But social categories that are orthogonal in the general population may be quite consolidated in a particular organization. For example, tenure and gender would be consolidated if the organization only recently began hiring women. And such consolidation is important because it reduces the degrees of freedom available to the manager in constructing the team. In such an organization, it is difficult simultaneously to increase the representation by women on a team and to increase representation by older cohorts. Moreover, in a situation in which fixed and variable parameters are consolidated-- say, where the human relations staff is all female and the sales department is all male-- if a project requires representation from both areas, the manager has limited discretion in determining its gender composition. Any assessment of the demography-based approach must therefore first clarify which demographic characteristics are fixed and which are variable and how these attributes interrelate in the organization's demographic makeup. An understanding of these issues can then guide analysis of the causal links depicted in figure 1 with a view to comparing the relative merits of the demography-based approach and an approach based on manipulation of the network-based social capital variables. We discuss next how we empirically validated our assertions about the limits to managerial discretion and the relative merits of demography-based and social networks-based approaches to staffing teams to achieve high performance.

METHOD

Our study population is "Malibu Research," a Midwestern contract research and development (R&D) firm that specializes in materials science and undertakes projects for two types of clients: its parent organization and outside firms in its local market. Projects fall into six basic types of services provided by Malibu: (1) scientific analyses, such as analyzing material properties or assessing product reliability; (2) conceptualization, or analyzing the potential and feasibility for new product ideas; (3) product/material development, which involves either developing a new product or assisting the client in developing a new product more efficiently; (4) process development, whereby Malibu helps the client to improve its process designs and flows; (5) manufacturing, which involves either performing material compounding for the client or assisting the client in manufacturing in-house or with a third firm; and (6) cost/quality initiatives, which include an array of services whereby Malibu assists clients in improving the cost or quality of a product after its initial launch.

We collected network and demographic data on 104 out of the 113 Malibu employees who worked on project teams during a one-year period following administration of a survey in the summer of 2001. We then obtained detailed data on the several hundred projects that were initiated by the firm in the year subsequent to our survey. The dataset has two unusual features that afford the testing of hypotheses that predict that network and other variables affect project team performance. First, since the network data were collected before the commencement of the projects whose performance we modeled, reverse causality can be ruled out. For instance, a team's internal network cannot become more dense because the members experience greater success at reaching their goals and objectives. Second, since each individual member of the firm participated in multiple projects, we are able to control for compositional differences between projects. In particular, we included a dummy variable for each employee who worked on at least two projects. This greatly reduces the problem of unobserved heterogeneity, whereby observed associations between networks and performance could be ascribed to differences in the ability or quality of team members. Rather, by including individual fixed effects, we attempted to account for variation in the extent to which a team is greater or less than the sum of its individual members. See Zuckerman, Reagans, and McEvily, (2004) for more detail on the advantage of the data set.

Following figure 1, we conducted three sets of analyses moving through the causal chain. First, we identified the dimensions of demographic diversity inside the organization that define the degree of latitude that a manager may have in managing a team's demography. Second, we considered how changes in a team's demographic diversity affect its social capital. And third, we modeled project performance as a function of teams' social capital, its demographic composition, and a series of control variables.

Dimensions of Managerial Discretion

To assess the range of discretion available to a would-be manager in designing a team, fixed characteristics that define the nature of the project must be distinguished from variable characteristics that are subject to manipulation by a manager charged with staffing the team. We assumed that the degree of functional diversity on a project reflects the underlying requirements for the task. For example, the service area "scientific analysis" draws on people from two functional areas, "materials" and "applied science." Since function is the fixed attribute, the variable attributes are other demographic qualities such as education, gender and tenure.

We analyzed the association between the functional areas and the variable attributes. Two notable attributes are missing from our analysis. The first is formal hierarchical position. Officially, Malibu is a flat organization. Every individual has the same title and reports directly to the chief executive officer. While there is no formal hierarchy in the firm, we suspect that an informal hierarchy does exist and that it is organized around knowledge. In knowledge-based organizations such as Malibu, more senior employees typically have more firm-specific knowledge, which enhances their standing in the informal hierarchy (e.g., Blau 1963). We also could not consider Malibu staff members' age since Malibu's human resources department does not record this information. This is a limitation of our study, since age is likely correlated with tenure, and age is the second most studied attribute in demography research (Tsui and Gutek, 1999: 36-37). There are also many other potential attributes that might be considered but that were not available to us.

We do, however, have data on most of the demographic characteristics that have been considered in past research. Table 1 is a cross-tabulation of function by each of these attributes for the Malibu staff members who participated in projects during the post-survey period. Each cell in table 1 contains a frequency count and a standardized residual (assuming independence between function and the focal attribute). Standardized residuals approximate a standard normal distribution with a mean of 0 and a variance of 1 and can be interpreted as Z-scores. Large Z-scores indicate significant over- and under-concentration in the firm. For example, women are more likely to work in business services than in any other area. There is also significant clustering by education and tenure. People with advanced degrees are more likely to work in applied science, while people who lack formal degrees are more likely to work in materials. With respect to tenure, individuals in analytical services are more likely to be employed by the firm for more than ten years, while an individual in materials is more likely to be a recent hire.

[TABLE 1 ABOUT HERE]

While table 1 reports bivariate associations between pairs of demographic categories, table 2 reports the results of a series of multivariate analyses among the categories. The F-statistics in column I are results from a discriminant analysis that tested for the various variable attributes' ability to discriminate among functional areas. The results indicate that there are significant differences between functions with respect to education, gender, race, and tenure. The strongest associations are for education, gender and race. The association between tenure and function, while significant, is weaker than the others.

[TABLE 2 ABOUT HERE]

The results in the last four columns of table 2 show the degree of consolidation among the variable attributes. The F-statistic in column II indicates a significant association between gender and education. The estimates in column III are based on a logistic regression of gender on the other variable attributes. These results show that the only significant predictor of being female is a high school diploma. The effect is negative and relative to the excluded category, which is a Bachelor's degree. That is, women are more likely than men to have earned a college diploma. None of the estimates in the last two columns are significant, indicating that both race and tenure are orthogonal to the other variable attributes, but for different reasons. Almost everyone is white, so being white is not salient. Not everyone has been with the firm for the same amount of time, but variation with respect to tenure does not vary greatly across the different functions and other background characteristics.

The results in table 1 and 2 indicate that the demographic landscape inside the firm has two primary dimensions: function and tenure. The other variable attributes-- education, race, and gender-- are essentially consolidated with function. Tenure is the second dimension, which has a weak association with function and is orthogonal to the other variable attributes. Thus, a manager's ability to staff teams based on demographic diversity is limited by the attribute landscape. Since function is the fixed attribute, the association between function and gender, education, and race limits a manager's discretion to manipulate the composition of a team with respect to these attributes. The manager has more degrees of freedom with respect to tenure. The significant variation across functions with respect to tenure means that a manager can manipulate the tenure distribution of a team while meeting the functional requirements of the task. Thus, our analysis provides a strong basis for focusing on function and tenure, which is largely consistent with the focus of past demography research (e.g., Ancona and Caldwell, 1992). More generally, the preceding analysis demonstrates the constraints that operate on the demography-based approach at Malibu. Because function and tenure represent the key dimensions of demography in the firm, diversity with respect to these attributes should influence the amount of social capital available to a team, as specified in figure 1.

Impact of Demographic Diversity on Teams' Social Capital

Social network variables. As detailed by Reagans and McEvily (2003), the network data were collected to examine how different features of network structure affect knowledge-sharing behavior. The data were collected using a combination of fixed-roster and free-recall approaches (Wasserman and Faust, 1994: 45-50). Each respondent was first presented with a fixed roster composed of a random sample of fifteen colleagues who had worked on the same projects as the respondent during the previous year. At Malibu, work is organized around project teams, so we suspected that network connections were shaped by such experiences. At the same time, two individuals could work on the same projects but never work with each other directly. Therefore, respondents were asked to cross out the names of colleagues with whom they had not worked directly. They were then instructed to copy up to ten of the remaining names to their contact list. Respondents copied a mean of nine names onto this list. Next, respondents were given two free-recall questions. They were first asked to list the names of colleagues who had been a significant source of knowledge during the previous year and then they were asked to list the names of colleagues for whom they had been a significant source of knowledge. In response to each name generator question, a respondent could list up to five colleagues, for a total of ten additional contacts. Respondents provided a mean of six additional names, with three unique names for each free-recall question.

The two procedures complement one another. The set of names generated by the free-recall questions supplements the fixed-roster, thus ensuring that important relationships that were not sampled for the fixed roster are added. The fixed-roster supplements the free-recall questions by ensuring that weaker, more formal ties are included in a respondent's network. Moreover, since we obtained the first and last names of each contact (even for those names generated by the free-recall questions), we could construct network matrices for the entire organization.

Despite its virtues, our data collection technique is not ideal. It would have been preferable to collect the network data using a complete sociometric roster. Unfortunately, our contact in the firm could not guarantee a high response rate if we used a sociometric roster because asking people to report on their relationships with over one hundred individuals is time consuming. Yet while falling short of the sociometric ideal, our network dataset compares favorably with past research in this vein. Previous studies on social networks and team performance have typically had complete information on the pattern of network connections inside a team but either lacked information on external connections altogether (Reagans and Zuckerman, 2001; Sparrowe et. al, 2001; Cummings and Cross, 2003;) or lacked information on how external contacts are connected to each other (Oh, Chung, and Labianca, 2003).

Despite the virtues of our data when compared with previous work in this area, we were concerned about the possibility that bias might be introduced by our data collection technique. There are three potential sources of bias that merit concern. First, our fixed roster put an artificial ceiling on the number of contacts that any one respondent could name (maximum of ten) and therefore might have prevented our observing important ties. Yet if respondents had more additional contacts, this would have been reflected in a tendency to provide more contacts in response to the free-recall questions. In fact, respondents named a mean of six out of a possible ten contacts for the free-recall questions. This suggests that, while our approach prevents respondents from naming as many contacts as would be possible with a standard sociometric instrument, most respondents would not have named any more than they did using our instrument, and the 20-contact limit would only have been problematic for respondents with a large number of contacts. To adjust for this possible truncation, we considered a tie to exist between two individuals if either individual cited the other one. When ties to or from a respondent were used to define respondents' contact networks, the size of such networks ranged from 3 to 49, with a mean of 22. Thus, while individuals with large networks may not have been able to name all of their contacts, such relationships would often be included through their contacts' responses.

A second source of potential bias concerns the distribution of tie strength. Under free-recall, respondents tend to cite their strong ties before their weak ties. As a result, if a survey limits the number of contacts a respondent can name, the sampled set of contacts could be skewed toward strong connections (Burt, 1986). Yet a careful examination of the data suggests that such a bias is unlikely to be large. In particular, our principal measure of tie strength was derived from respondents' answers to the question "on average, how often do you talk to [each contact]?" Following Burt (1992: 125), we also considered an alternative measure of tie strength, based on respondents rating of "how close (they) are with each [contact]" on a scale from "especially close," through "close," "not so close," and "distant." The association between emotional closeness and communication frequency is strong at Malibu. Given the strong association between the two tie-strength measures, we calculated our network variables using communication frequency as the indicator of tie strength. Results were the same when we constructed those measures using emotional closeness to define tie strength or when we defined tie strength as the mean or product of emotional closeness and communication frequency. Across all respondents, 39 percent of the contacts are spoken to "daily", 34 percent are "weekly", 16 percent are "monthly", and 10 percent are "less often." That there is considerable variation with respect to tie strength suggests that, while our instrument may have truncated the set of contacts that a respondent might have mentioned in a standard sociometric survey, the truncation was not so severe that respondents only had space for strong ties but not for their weak ties.

A third potential source of bias is the most serious: that our network variables would not be reliable indicators of the true amount of social capital available to a team. We first discuss how the network variables were constructed and then return to the potential bias contained in these measures.

We measured internal density as the mean strength of connections between project members:

$$ND_k = \sum_{i=1}^{Nk} \sum_{j=1}^{Nk} \frac{z_{ij}}{\max(z_{iq})} \Big/ N_k(N_k - 1),$$

where z_{ij} is the strength of the tie from team member i to member j, $\max(z_{iq})$ is the strongest of i's reported ties to anyone in the firm, N_k is the number of members in team k, $N_k(N_k - 1)$ is the maximum number of ties among members of team k. Scaling by $\max(z_{iq})$ removes individual differences in the tendency to report high tie strength to others. Network density varies from 0, no relations between team members, to 1, maximum strength relations between all team members.

Our measure of external range is the reverse of Burt's (1992) measure of network constraint:

$$ER_k = 1 - C_k$$

where the constraint posed by alter j on ego i is measured as in Burt (1992: chap. 2) and averaged over all team members:

$$C_k = \sum_{i=1}^{Nk} \sum_{j=1}^{Ne} (p_{ij} + \sum_{q=1}^{N} p_{iq} p_{qj})^2 \Big/ N_k;$$

where N_e is the number of contacts external to the team.

Constraint has two components. The first is the proportion of his or her total relational strength that i directly allocates to j:

$$p_{ij} = (z_{ij} + z_{ji}) \Big/ \sum_{q=1}^{N} (z_{iq} + z_{qi})$$

The second component is the strength of the indirect connection between i and j through mutual contacts q, where p_{iq} is the proportion of his total relational strength that i devotes to q and p_{qj} is the proportion of her total relational strength that contact j devotes to contact q:

$$\sum_{q=1}^{N} p_{iq} p_{qj}$$

The more that team members are connected to the same contacts q both inside and outside the team, the more they are redundant with one other. The higher the teams' mean level of constraint, the lower the team's external range. We focus on the reverse of constraint rather than constraint itself simply so that we can frame the discussion in terms of how the network variables increase rather than decrease performance. And we describe this as a range measure because constraint itself is a variant of the family of range measures described by Burt in earlier work (see Burt, 1983). In a companion paper (Reagans et al., 2004), we provide further discussion of the theoretical issues that underlie this measure and provide alternative measures of constraint and external range that are even better suited for measuring breadth of access to information and resources. The results based on this alternative strengthen the findings presented below.

We now return to the possibility that our social network instrument may not have produced reliable estimates of Malibu teams' levels on internal density and external range and thus the true amount of social capital available to the team. One way to assess this issue is to make a direct comparison between the level of internal density calculated from similar sociometric data and the level of internal density in our study. The network data from Reagans and Zuckerman (2001) were collected using a complete team roster. Their teams were similar to ours in several respects. The teams were about the same size, and demographic diversity was defined by firm tenure. The level of communication between team members was measured using almost identical response categories. At the same time, while the teams studied by Reagans and Zuckerman typically had more long-term research objectives, such as patenting, the Malibu projects were contract research and development work and had more short-term goals.

Because the teams studied by Reagans and Zuckerman were longer-lived and the network data from those teams were collected using a sociometric roster, we expected internal density to be higher in their study. Yet the difference turned out to be slight. In particular, the log of internal density had a mean of .43 (with a standard deviation of .13) in the Reagans and Zuckerman (2001) data, as compared with a mean of .38 (with a standard deviation of .07) in our study. The difference in internal density across the two studies was .05, and our mean level of internal density was within one standard deviation of their mean. We also checked to see if the difference across the two studies increased with team size. If our data collection technique excluded network connections, then the size of the bias it introduced should have gotten bigger as team size grew. Therefore, one would expect the difference in internal density across the two studies to increase with team size. For each category of team size, we calculated the average internal density score and performed a paired t-test for internal density, with team size as our grouping variable. There was a small tendency for the teams studied by Reagans and Zuckerman to be more internally dense, but the difference (.02) was small and was not statistically significant. In addition, we found that the difference actually got smaller as team size increased. When team size was greater than or equal to 20, our teams are more internally dense than the teams studied by Reagans and Zuckerman. These results seem to suggest that whatever bias was introduced by our method is small and more importantly, it does not appear that the size of the bias increases with team size. We cannot make a similar comparison for external range because Reagans and Zuckerman did not have data for networks beyond the team.

Demographic diversity. Functional diversity was defined as:

$$DV_k = 1 - \sum_{c=1}^{C} P_c^2$$

where C is the number of functional areas, and P_c is the proportion of team members from area c (Blau, 1977; Allison, 1978). The measure is a function of the distribution of people across different functional areas. For example, demographic diversity with respect to function would be high if a team is composed of people from a large number of functions with equal representation for each area. We measured diversity on organizational tenure as the coefficient of mean difference (CMD) in tenure. The variable measures the average difference in tenure between project members,

$$CMD_k = \frac{1}{N_k(N_k - 1)} \sum_{i=1}^{N_k} \sum_{j=1}^{N_k} |t_i - t_j|, \; j \neq i$$

where t_i and t_j are the tenure of person i and person j, respectively (Kendall and Stuart, 1977: 48). Often demographic diversity is measured with the Gini index (Gini) or its near equivalent, the coefficient of variation (CV). The coefficient of mean difference differs from these measures in that the Gini and the CV are divided by twice the group mean. The Gini and the CV are scaled by the mean to capture the intuition that holding the dispersion of a resource constant, an increase in the average level of that resource lowers the felt level of inequality (Allison, 1978: 467). While this transformation often makes sense for an inequality measure, it is not clear if the transformation is appropriate for an indicator of demographic diversity. With respect to organizational tenure, for example, this would imply that for two equally diverse projects, the project with higher mean tenure would be less diverse than the project with lower mean tenure. We avoid having to make this troubling assumption by controlling for mean tenure in the models we present rather than dividing our diversity variables by their means (Sørensen, 1999, 2000; cf. Reagans and Zuckerman, 2001).

Control Variables. We also included three control variables in our analysis. First, since the density of a network tends to be a negative function of size, we controlled for the number of people who had ever worked on the project (team size). Second, we controlled for the extent to which project members worked with one another in the prior year (team members' shared prior-year experience). While our analysis focused on projects that started after the network survey, projects prior to the network survey were used to define the sampling frame for formal contacts. Thus, it is important to identify that aspect of network structure that is due solely to the construction of the survey versus that which reflects "true" network patterns. To that end, we used detailed data on the number of hours each person allocated to a project during the previous year to calculate the tendency for each pair of project members in the current year to have worked together in the prior year:

$$PH_k = \sum_{i=1}^{Nk} \sum_{j=1}^{Nk} (s_i + s_j)/(total_i + total_j) \Big/ N_k (N_k - 1)$$

where s_i is the number of hours that person *i* allocated to projects that included j, s_j is the number of hours that person j billed to the same projects, and $total_i$ and $total_j$ are total hours they each billed in the previous year. It is important to note that our control for shared prior year experience does not necessarily guarantee that team members interacted with each other in the previous year. For instance, team members may have worked on the same project at different points in time. Consequently, this control variable is a proxy for shared prior year experience that may include some measurement error. Third, we also controlled for the average number of hours that members of the project allocated to projects that were ongoing (team members' shared concurrent experience*)*. The pattern of network ties could be a function of projects that started in the previous year but continued into the current year. Using the same billing data, we calculated the tendency for project members to overlap in their billable hours on ongoing projects. The equation is identical to the one used to measure hours allocated to projects from the previous year, except that we considered only on-going projects.

Analysis. Descriptive statistics and a correlation matrix for the variables described above are presented in table 3. External range and internal density have a positive, non-negative, and significant correlation (r = .383), which is consistent with our argument that the two forms of social capital are not in tension with one another. The positive association suggests differences among teams in terms of their members' overall sociability, or general propensity to develop network connections across the organization, with a greater propensity increasing both internal density and external range.

The statistics and those in the analysis in table 4 are based on 10,544 observations. The number of observations is large because team characteristics vary over time, so a single project is observed multiple times. To control for non-independence, we introduced N-1 project dummies (i.e., fixed effects) into the regression equation describing the association between demography and the network variables (Hannan and Young, 1977). Because we included fixed effects, the estimates should be interpreted as explaining within-team variation. That is, we estimated the impact on the density and range of a team's network when a new member was added and consequently brought about change in the demographic diversity and/or any of the control variables.

[TABLE 3 ABOUT HERE]

There are two equations in table 4, one in which the dependent variable is the natural logarithm of internal density and the other, the natural logarithm of external range. We logged both variables due to high skewness in each. Owing to this transformation, the effects should be interpreted as displaying percentage increases or decreases in either internal density or external range that is caused by a one-unit change in the predictor. The estimates for team size show that initial increases in team size have a positive effect on internal density, but when team size equals approximately ten members, continued increases in team size start to reduce internal density. Since the coefficients are within-project estimates, the pattern most likely reflects the tendency for people who communicated with each other more often at the start of the period to join projects at proximate points in time and for people with weaker connections to these individuals to join later. In addition, as would be expected, the tendency for project members to have worked together in the prior year strongly increases the project's density. Perhaps surprisingly, teams whose members worked often with one another in the prior year also tended to exhibit greater external range.

[TABLE 4 ABOUT HERE]

These models also provide support for the hypothesized direction of both β_1 and β_2 from figure 1. Demographic diversity, defined either by function or tenure, has a significant and negative effect on internal density. Diversity of either kind makes it less likely that team members are connected by strong network connections. But, as predicted, increasing diversity is also associated with more external range. The members of diverse teams are more likely to be connected to different external contacts, and those contacts are more likely to be disconnected from each other. When compared with tenure-based diversity, functional diversity appears to have a much larger effect on the network variables. This is unsurprising because the most influential force in organizing the informal social network of an organization is its formal organizational network (e.g., Han, 1996). Thus we have shown that demographic diversity has the effects on team social capital that are predicted on the left side of figure 1 and that these effects are larger for diversity defined by function than for diversity defined by tenure. The raw coefficients also seem to suggest that the decrease in internal density caused by a rise in demographic diversity is greater than the increase in external range that is caused by such a rise. These differences are largely a reflection of the greater variance in the internal density measure. Accordingly, the standardized coefficients for tenure diversity are of essentially the same absolute magnitude in both models: -.040 in the first equation and .041 in the second. The absolute value of the standardized effect of functional diversity on logged internal density (-.191) is three times greater than the comparable effect on logged external range (.068).

The network at Malibu research is homophilous with respect to tenure and function, which produces the expected effects of demographic diversity on both internal density and external range.

Analyzing Project Performance

Next, we attempted to validate the right half of figure 1, which predicted a positive association between each of the social capital variables and team performance. We use project duration as our performance measure (cf. Hansen, 1999). In particular, we modeled the hazard of a project's completion on a given day after it was first initiated. Because the firm bills clients on a cost-plus basis, longer-lived projects may generate more revenue because they typically involve more labor input. Yet the firm's management informed us that they regarded the duration of a project, conditional on the amount of labor invested and the nature of the task (for which we included control variables), to be a useful measure of poor performance because there are strong incentives for such hours to be concentrated in a minimum of calendar days. In particular, delay hurts the firm's reputation with a particular client, which reduces the probability of a repeat engagement. An individual at Malibu aptly summarized this view when he stated "If you overstay your welcome, there's a perception that you don't deliver. Everybody wants it done yesterday. Deliver quickly and they will always invite you back." We thus modeled the hazard of project completion as a function of our network, demographic, and other measures while controlling for labor input.

Modeling framework and controls. Duration was analyzed using a continuous-time event history model. Project duration is the number of days between the day the first project member billed hours to the project and the day that the customer was billed, though we considered whether there is delay between when a team completed a project and when the firm billed the customer. The natural billing cycle of the firm could add noise to project duration, but it turns out that the distribution of billing days was not concentrated on any one particular day: 1.2 percent of the bills occurred on Sunday, 27.8 percent on Monday, 19.8 percent on Tuesday, 18.8 percent on Wednesday, 17.1 percent on Thursday, 12.1 percent on Friday and 3.2 percent on Saturday. That a higher percentage of projects was completed on Monday most likely reflects a tendency for products that were completed on either Saturday or Sunday to be billed the following Monday. This suggests that any billing cycle effect at most adds one or two days to the duration variable. A second concern is the potential difference between projects that bill a customer multiple times and projects that billed a customer once or not at all. Success with a particular customer could be reflected in the tendency for a team to bill the customer multiple times, so it was important for our teams to be at risk for multiple completions. Each time a team billed a customer, project duration was reset to zero, and the team was at risk for another completion. In addition to allowing for multiple completions, we controlled for the <u>number of past completions</u>. This allowed us to focus on the factors that speed project duration and therefore free a team up for another completion, while controlling for the number of past completions.

This modeling approach has several appealing features. First, it allows for multiple project completions. As described above, this is important for distinguishing more successful from less successful projects. Second, teams that never billed a client (e.g., censored observations) contribute to the parameter estimates. This is important because it rules out the possibility that our parameter estimates are biased because we sampled on completed projects. Third, the analysis is dynamic. The predictors in the model co-vary with time. We assume that a person joined a project on the first day he or she billed hours to the project. The project variables were updated each time a new person joined a project or when there was a change in any of the independent variables. Therefore, we could examine how the change in each predictor speeded project completion.

Of the 1,518 projects, 785 were completed by the closing of our observation window. Table 5 presents results from a continuous-time hazard model estimated using the "streg" procedure with the stata statistical package. The error term in the model is assumed to have a log-logistic distribution. Our substantive conclusions are robust to this assumption. For example, we get the same pattern of results if we assume a weibull or log-normal distribution. In addition to adjusting for multiple observations from the same project, we used robust standard errors. Further, as discussed above, the models include dummies for each person who participated in multiple projects. This strategy helps us avoid two sources of unobserved heterogeneity. First, the inclusion of individual dummies accounted for any differences in team performance that merely reflect the tendency of certain projects to attract more high-performing individuals than others. In addition, the individual dummies are an effective means of controlling for differences in project type. In general, projects of different types should be distinct in the sets of employees who generally work on them.

Thus, including the individual dummies allows the duration of a given project to be evaluated in relation to a baseline expectation derived from the nature and the quality of the labor input to the project. As discussed below, functional heterogeneity is also a useful control in this regard because it should reflect the complexity of the project. Controlling for the number of labor hours allocated to a project is also important because this variable indicates the scale of the project and could indicate the project's priority. More important projects attract more labor-hours that could be billed to other projects. We thus included in all models a control for the cumulative number of labor-hours billed to a project:

$H_k = \sum_{i=1}^{Nk} h_i$,

where h_i is the total number of hours that person i allocates to project k (cumulative labor-hours devoted to project). Finally, we controlled for team size, which we measured as a count of the number of people who had billed at least one hour to a project by day t.

Finally, we ran additional models beyond those presented in table 5 to address unobserved heterogeneity. First, we estimated a model that included six additional variables, each of which registered the number of people from each function that were represented on a project. Together, these variables measure differences between projects as a matter of the extent to which the projects combined elements of different functions in varying proportions. Second, we estimated "frailty" models that represent unobserved heterogeneity as a random variable that influences the hazard rate as a disturbance term. All the effects that are significant in model 5 were significant in a comparable frailty model as well.

[INSERT TABLE 5 ABOUT HERE]

Analysis. The first model in table 5 contains the individual-level dummy variables that constitute our main method of controlling for unobserved heterogeneity. The next model adds controls for project scale and members' concentration of their work with one another. The results indicate that project duration increases with the number of past completions. Perhaps with each additional completion, the task becomes more challenging and demanding, which increases the duration of the next assignment. The results also show that concentration of labor hours either on past projects or ongoing projects delays completion time. Perhaps the concentration of labor hours outside the focal project delays the completion of the focal project because project members work on multiple activities at once, making it more difficult for them to accomplish stated goals and objectives. It is not clear if the knowledge accumulated while working together the previous year made members of the focal project more complacent and therefore less productive, or if the knowledge project members accumulated while working together in the previous year degrades their performance in the current year.

The effects of demographic composition are introduced in the second model. Neither of the diversity measures has a significant impact on project duration. If we do not include the number of past completions in our model, however, functional diversity has a marginally significant positive effect. This result suggests that projects that draw on a wide variety of functions are likely to be more complex and challenging than those that are restricted to members of a single function. As such, the inclusion of functional diversity serves as an additional corrective for unobserved heterogeneity beyond the individual dummies and project scale. There is no discernible effect for tenure diversity, which is the primary characteristic available for manipulation by Malibu managers. To recall, our analyses in tables 1 and 2 suggest that the other variable characteristics—education, gender, and race--are either consolidated with function or display too little variation to be managed effectively. We also computed team diversity measures with these variables, but none of them have significant associations with project duration.

Lawrence (1997) provided one reason why a team's demographic diversity may have no impact on its performance: the demographic characteristics in question may have a weak association with the organization's social networks or other factors that mediate between demography and performance. Yet as we demonstrated in table 4, the Malibu social network is significantly organized by both function and tenure. In particular, diverse teams are generally less dense and exhibit more range. Moreover, as shown in models 3 and 4 of table 5, each of the network-based social capital variables that is thought to mediate between demographic diversity and performance is significant and in the direction postulated in figure 1. The estimate for internal density is negative and significant. The estimate for external range is also negative and significant. Since faster completion times indicate better performance, both network variables have a positive effect on team performance. The observed pattern of effects is consistent with our argument that the network variables are the more proximate determinants of performance. The results also reinforce the notion that the optimal network structure for a team is characterized both by high internal density and high external range. Local structural holes limit team performance, but global structural holes improve performance.

To gain greater insight into the substantive significance of the effects of the social network variables on project duration, in figure 3 we plotted the estimated survivor function based on model 5 in table 5 and at three levels of the two network variables. The largest estimated differences occur from day 30 to day 50. Conditional on having survived to day 40, for instance, a project with mean internal density and external range is expected to end with a probability of about .62 (1-.38). But increases in both external range and internal density of one standard deviation elevate this probability to about .74 (1-.26), and decreases by that amount lower the probability to about .45 (1-.55). These patterns suggest substantial differences in the speed at which projects are completed based on variation in their social networks.

[FIGURE 3 ABOUT HERE]

Thus, just as the results presented in table 4 provide support for the left half of figure 1, the results in table 5 support the right half. Then why, if each of the causal pathways in figure 1 has been validated, is manipulating demographic diversity irrelevant to the performance of these teams, as suggested by the results in table 5? As we argued above, the reason derives from the ambiguous implications of demographic diversity, even when it affects network structure in the manner expected by theory. Since, under conditions of homophily, a manager who increases diversity will both increase external range and decrease internal density, and since both external range and internal density have positive effects on performance, the overall implications of managing demographic diversity are unclear. Moreover, as we have seen, the demographic makeup of Malibu narrowly restricts the degree of discretion for managing team demography, especially with respect to variables that are salient in the organizational environment (e.g., race, education, gender). The only form of demographic diversity that has even a marginal overall association (δ in figure 1) is functional diversity, and it only has an effect when we do not include the number of previous completions in the model. Moreover, functional diversity is essentially a fixed characteristic that defines the nature of the project rather than a variable that is subject to managerial discretion.

DISCUSSION

Difficulties with the Demography-based Approach

We found considerable evidence in support of the framework describing how demographic diversity affects team performance (figure 1). At Malibu, increasing demographic diversity inhibits internal density and promotes external range. And each of these network variables has a positive effect on team performance. Such evidence might seem to support the assertion that the productivity of a team can be managed by manipulating its demographic make-up so as to alter the intervening network variables. But closer inspection of both past research and the argument and evidence presented here gives reason to doubt the effectiveness of attempting to manage team demography. The first reason for such skepticism was articulated by Lawrence (1997), who argued that the assumption of strong homophily, whereby the demographic variables essentially determine network structure, does not hold in most organizational settings.

The present paper has presented two additional reasons why a demography-based approach to managing diversity is a problematic practice for managers. The most basic issue concerns the range of discretion that is available to vary the demographic composition of project teams. The association among demographic characteristics in terms of their tendency to co-occur as attributes of the same organizational members defines the demographic make-up of an organization. This landscape can differ between organizations and within the same organization across successive waves of cohorts, depending on the demographic profile of the people being hired and those who leave (see Sørensen, 2003). For example, a woman who joins a team at Malibu Research is more likely than not to work in business services. We found that function and tenure effectively define the attribute space at Malibu. People in the organization also differ with respect to education, gender and race, but these attributes are highly associated with function. And the amount of functional diversity on a project represents the nature of the task being performed. So while education, gender, and race are in theory open to manipulation, in practice, they are largely fixed. Knowledge of the demographic make-up of an organization helps a manager avoid predictable errors. For example, a manager could attempt to maximize (or minimize) diversity with respect to an attribute that is correlated with the fixed attribute. It would also be problematic to attempt to maximize (or minimize) diversity with respect to two variable attributes that are negatively correlated. The same issues also pertain to research that tends to treat demographic variables as independent of one another. The correlation among different demographic characteristics can produce mixed results across empirical studies.

The second and deeper reason for challenging a demography-based approach is that a stronger correlation between team demography and the network variables paradoxically undermines, rather than supports, such a rationale. Under such conditions of significant homophily, increasing diversity tends both to reduce internal density and to heighten external range, and both network variables tend to facilitate higher team performance. Manipulating diversity thus produces a trade-off between the two causal pathways. The size of the trade-off depends on the relative strength of the paths in the causal chain. In some contexts, the top pathway will be stronger than the bottom pathway, thus producing a negative association between diversity and performance. In other situations, the reverse will be true. In still others, the two causal pathways are balanced such that there is no significant association, as we found at Malibu Research. The fundamental difficulty with both research and policies that focus on the effect of demographic diversity on performance derives from the difficulty of predicting which pathway will be stronger than the other, because there is little theoretical or empirical basis for the prediction. This thereby provides a poor foundation for predicting the impact of increasing diversity even when support is found for the full causal framework.

How Good is the Network Alternative?

The two issues we have raised with a demography-based approach to team evaluation and design suggest that the alternative, social network-based approach might be preferable, at least in certain instances. The effectiveness of this strategy clearly depends on how well the manager knows the organization's informal network. Krackhardt (1990) has shown that managers often do not have a good sense of the network structure of the organization, particularly of the ties among employees that are several levels below managerial rank. So the difficulties associated with managing by team demography have to be weighed against the costs associated with collecting network data necessary to obtain an accurate view of networks.

Until recently it has been very difficult to collect network data. Moreover, while the growth and increasing sophistication of network surveys, annual peer reviews, e-mail logs, and other archival data make it easier to collect network data, the costs of collecting such data remain significant. For example, the network data reported here are based on the responses of over 100 individuals, who on average took one-half hour to complete the survey. If we add this time to the number of hours it required to design and administer the survey, as well as compute the basic analyses necessary for creating the network measures, we estimate that 200 man-hours were expended.

While these labor costs are not trivial, it appears that they are small relative to the value of the information provided by the network data. In figure 4, we provide an illustration that gives a rough sense of the potential savings based on the estimated survival plot presented in figure 3. To construct the graphs in figure 4, we considered the 1,518 projects initiated by Malibu during our observation window under two conditions: if the teams that undertook these projects were at the mean levels of both internal density and external range, and if they were one standard deviation above the mean on these variables. As depicted in figure 3, the estimated survival rate of each project in the second set of projects would be lower, especially from days 30 to 50. Using the survival curves in figure 3, we estimate the number of projects that would have been completed on a particular project day, under both network conditions. By taking the difference between these two numbers, we have an estimate of the additional number of projects that would have been completed if both network variables were one standard deviation above the mean. The circles in figure 4 are the predicted number of additional projects that would have been completed. Finally, if these extra project completions are multiplied by the mean number of hours billed by projects of that age, we derive the number of labor hours saved.

[FIGURE 4 ABOUT HERE]

The results in figure 4 suggest that the benefits of using the network variables far outweigh the costs. For example, Malibu is estimated to save over 2,000 labor hours on day 17 of the projects. Admittedly, it would be difficult for a manager to assign individuals to projects in which every project is at the mean or even one standard deviation above the mean. Another interpretation for the figure, however, is the additional number of labor hours that would be saved if some fixed proportion of the projects were designed using the network criteria. The savings for these projects would start on day one and continue until they were completed. For example, if a manager used the network criteria and constructed 30 percent of the total projects such that each one was one standard deviation above the mean with respect to internal density and external range, those projects would save the firm more than 2,200 labor hours by day 17. Those savings are eleven times the estimated cost of administering the network survey and only require that the manager use the network criteria to manage a fraction of the total projects.

Overall, the evidence indicates that although it may seem easier to manage team performance using demographic criteria, it is more effective to use network criteria. The network variables are more precise and proximate indicators of the causal mechanisms than are demography variables. And the overall effect of network variables is more predictive of team performance than demography variables. Consequently, using the network criteria to build teams would improve their performance in this organization. We believe the same is true generally, because of the inherent challenges associated with improving team performance by manipulating its demography.

Limitations and Future Directions

We regard our results as having strong implications for the relative merits of the demography-based approach and a social-network-based alternative. Yet we must also highlight several limitations with our analysis. First, while our analysis includes important demographic characteristics such as functional area, tenure, race and gender, it does not include age or hierarchical position, variables that have been show to have an impact on important organizational outcomes. Second, our results come from a single organization, so it is not clear how well our results generalize to other organizations that differ on key dimensions. For instance, it is possible that Malibu's flat organizational structure may act to strengthen the association between team demography and the network variables and between the network variables and team performance. We suspect that similar network dynamics occur in other organizations, with respect to the origin of the network variables and their impact on team performance, but we await future research to confirm our conjecture. Third, as noted above, our social network data fall short of the ideal complete, sociometric data on the entire organization. While we think our method is quite useful in a context in which it is difficult to use a sociometric instrument, and our analyses suggest minimal bias, we should not lose sight of the ideal. Future studies should endeavor to use sociometric instruments wherever possible. In addition, future research would do well to consider multiple performance metrics. While our fieldwork makes us highly confident that project duration is a useful measure of performance at Malibu, a richer picture might have emerged were we to have additional performance metrics.

Finally, it is useful to consider the implications that flow from a comparison of our framework with what might be termed the "moderator approach," due to its focus on factors that can reduce the amount of homophily inside a collective unit (e.g., a team or an organization) (e.g., Chatman et al., 1998; Wesphal and Milton, 2000; Ely and Thomas, 2001; Polzer, Milton, and Swann, 2002). By reducing the amount of homophily inside the unit, these factors limit the negative effect that diversity has on internal density, thus allowing the positive effect on external range to dominate. The prospect of moderation provides a manager with an alternative way to staff such a unit when drawing members from a larger environment (i.e., the organization, if the unit is the team, or the wider society or market, if the unit is the organization) that is characterized by significant homophily. First, the manager maximizes the team's demographic diversity. Given the significant homophily in the environment, the team should enjoy high external range, but it should also suffer from a low level of internal density. But the manager need not settle for this low level of internal density. Rather, he or she may be able to build greater internal density, and thus a greater capacity for coordination, through processes and team-building exercises that help members of the different demographic categories to see their different perspectives as sources of distinct value rather than bases for opposing identities (Ely and Thomas, 2001). By offsetting the negative effect that demographic diversity has on internal density but leaving intact the positive effect on external range, these managerial interventions allow the promise of diversity to be realized.

While we agree that this is a viable alternative in some situations, we also think that the moderator approach makes a critical assumption that limits its scope. The approach assumes that contextual factors reduce the level of homophily within the collective unit in question but not in the larger environment within which the unit operates. If the collective unit is a single team or a small number of teams, this assumption seems plausible. But if moderation happens at the firm level, the assumption is more problematic. The causal structure depicted in figure 1 implies that such a reduction in homophily would render both the top and bottom causal pathways insignificant. Since the degree of homophily inside the firm governs both network patterns, reducing the negative association between diversity and internal density also serves to offset the positive association between diversity and range inside the organization. Therefore, in order for diversity to have a positive effect on performance, there must be a positive association between demographic diversity and external range in the network <u>outside</u> the firm. And even if there is such an association, it is not clear that external range outside the firm is often important for performance. For instance, while it is clearly useful for product development teams to be ethnically diverse if the product it is developing is to be tailored for different ethnic communities, there are many purposes for which the expansion in external range provided by such diversity holds little value in furthering a team's objective. More generally, increasing external range outside an organization does not help a team meet the crucial challenge we identified at the outset of this article: to access a wide range of the information, resources, or perspectives that are distributed throughout an organization.

Furthermore, while the moderator approach focuses on instances in which relationships are forged across demographic divides within an organization or team, it actually assumes that such divides remain important bases for homophily in the environment outside such social units. And just as it is problematic to assume that networks and demography are equivalent within organizations (Lawrence, 1997), such an assumption may also not hold in the environment for the particular demographic characteristic under consideration. As such, there may be value in using the social networks based approach we have developed here to apply a version of the moderator approach that eschews the assumption that external networks are characterized by homophily. For instance, rather than maximize demographic diversity in an indirect effort to maximize external network range and then counteract the potential for low internal density by helping people overcome their demographic differences, a firm could directly maximize external range and internal density, as well as develop such density after the initial formation of the unit.

Finally, this discussion points to the tradeoffs associated with our strategy of collecting social network data during the period prior to the initiation of the projects whose performance we model. On the one hand, this approach is very attractive as a means for addressing the issue of reverse causality, which tends to plague research on the effect of social networks on performance (see Zuckerman, Reagans, and McEvily, 2004 for a more extended discussion). On the other hand, our strategy requires us to assume that the network in one year was largely determined by the network in the previous year. A more complete analysis would include data from both years. More generally, future research should endeavor to examine directly the social processes that occur on a team and help mediate the effect of social network variables on performance, rather than having to treat them as a "black box," as we were forced to do in this analysis.

Deciding who to put on a project or a team is one of the biggest challenges facing a manager or team leader. Relying on demographic criteria has been favored, both on practical and theoretical grounds. The argument and evidence presented here challenges the use of demographic criteria on both grounds. First, there is little theoretical basis for predicting how changes in the demographic composition of a team will affect team performance.

Second, even if there were, the demographic make-up of a firm quite often makes it difficult for a manager to manipulate a team's demographic composition successfully. Our results suggest that relying on social networks to guide staffing decisions would have clearer performance implications. We therefore conclude that a social network-based approach is preferable to an exclusive focus on team demography. We recognize, however, that a deeper understanding of the value of the approach we have proposed would come from implementation by a manager or, at the very least, via a more developed simulation. Such an exercise would likely provide additional insight into the opportunities and limitations that are afforded by a focus on the social network variables.

REFERENCES

Allison, P. D.

1978 "Measures of inequality." American Sociological Review, 43: 865-880.

Aldrich, H. E., and C. Zimmer

1986 "Entrepreneurship through social networks." Pp. 3-24 In D. L. Sexton and R. W. Smilor (eds.), The Art and Science of Entrepreneurship: 3-24. Cambridge, MA.: Ballinger.

Ancona, D. G., and D. F. Caldwell

1992 "Demography and design: Predictors of new product team productivity." Organization Science, 3: 321-341.

Bantel, K. A., and S. E. Jackson

1989 "Top management and innovations in banking: Does the composition of the top team make a difference?" Strategic Management Journal, 10: 107-124.

Blau, P. M.

1963 Dynamics of Bureaucracy: A Study of Interpersonal Relations in Two Government Agencies. Chicago: University of Chicago Press.

1977 Inequality and Heterogeneity. New York: Free Press.

Blau, P. M., and J. E. Schwartz

1984 Crosscutting Social Circles: Testing a Macrostructural Theory of Intergroup Relations. Orlando, FL: Academic Press.

Brass, D. J.

1985 "Men's and women's networks: A study of interaction patterns and influence in an organization." Academy of Management Journal, 28: 327-343.

Brewer, M.

1979 "Ingroup bias in the minimal intergroup situation: A cognitive-motivational analysis." Psychological Bulletin, 86: 307-324.

Burt, R. S.

1980 "Autonomy in a social typology." American Journal of Sociology, 85: 892-925.

1982 Toward a Structural Theory of Action: Network Models of Social Structure, Perception and Action. New York: Academic Press.

1983 "Range." In R.S. Burt and M.J. Minor (eds.), Applied Network Analysis: 176-194. New York: Academic Press.

1986 "A note on sociometric order in the General Social Survey network data." Social Networks, 8: 149-174.

1992 Structural Holes: The Social Structure of Competition. Cambridge, MA: Harvard University Press.

2000 "The network structure of social capital." In B. M. Staw and R. I. Sutton (eds.), Research in Organizational Behavior, 22: 345-423. Greenwich, CT: JAI Press.

Byrne, D. E.

1971 The Attraction Paradigm. New York: Academic Press.

Chatman, J. A., J. T. Polzer, S. G. Barsade, and M. A. Neale

1998 "Being different yet feeling similar: The influence of demographic composition and organizational culture on work processes and outcomes." Administrative Science Quarterly, 43: 749-780.

Coleman, J. S.

1988 "Social capital in the creation of human capital." American Journal of Sociology, 94: S95-S120.

Cummings, J. N., and R. Cross

2003 "Structural properties of work groups and their consequences for performance." Social Networks, 25: 197-210.

Ely, R. J.

1994 "The effects of organizational demographics and social identity on relationships among professional women." Administrative Science Quarterly, 39: 203-238.

Ely, R. J., and D. A. Thomas

2001 "Cultural diversity at work: The effects of diversity perspectives on work group processes and outcomes." Administrative Science Quarterly, 46: 229-273.

Feld, S.

1982 "Social structural determinants of similarity among associates." American Sociological Review, 47: 797-801.

Gabbay, S. M., and E. W. Zuckerman

1998 "Social capital and opportunity in corporate R&D: The contingent effect of contact density on mobility expectations." Social Science Research, 27: 189-217.

Granovetter, M. S.

1973 "The strength of weak ties." American Journal of Sociology, 78: 1360-1380.

Hackman, J. R., and C. G. Morris

1974 "Group tasks, group interaction process, and group performance effectiveness: A review and proposed integration." In L. Berkowitz (ed.), Advances in Experimental Social Psychology, 8: 45-99. New York: Academic Press.

Han, S.

1996 "Structuring relations in on-the-job networks." Social Networks, 18: 47-67.

Hannan, M. T., and A. A. Young

1977 "Estimation in panel models: Results on pooling cross-sections and time series." In D. Heise (ed.), Sociological Methodology 1977: 53-87. San Francisco: Jossey-Bass.

Hansen, M. T.

1999 "The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits." Administrative Science Quarterly, 44: 82-111.

Homans, G. C.

1950 The Human Group. New York: Harcourt, Brace.

Ingram, P., and P. Roberts

2000 "Friendships among competitors in the Sydney hotel industry." American Journal of Sociology, 106: 387-423.

Kendall, M., and A. Stuart

1977 The Advanced Theory of Statistics, vol. 1, 4th ed. New York: Macmillan.

Krackhardt, D.

1990 "Assessing the political landscape: Structure, cognition, and power in organizations." Administrative Science Quarterly, 35: 342-369.

Kramer, R. M.

1991 "Intergroup relations and organizational dilemmas: The role of categorization processes." In L.L. Cummings and B.M. Staw (eds.), Research in Organizational Behavior, 13: 191-228. Greenwich, CT: JAI Press.

Laumann, E. O.

1966 Prestige and Association in an Urban Community. New York: Bobbs-Merrill.

Lawrence, B. S.

1997 "The black box of organizational demography." Organization Science, 8: 1-22.

March, J. G.

1991 "Exploration and exploitation in organizational learning." Organization Science, 1: 71-87.

McCain, B. E., C. O'Reilly, and J. Pfeffer

1983 "The effects of departmental demography on turnover: The case of a university." Academy of Management Journal, 26: 626-641.

McPherson, M., L. Smith-Lovin, and J. Cook

2001 "Birds of a feather: Homophily in social networks." Annual Review of Sociology, 27: 415-445.

Mehra, A., M. Kilduff, and D. J. Brass

1998 "At the margins: A distinctiveness approach to the social identity and social networks of underrepresented groups." Academy of Management Journal, 41: 441-452.

Oh, H., M. Chung, and G. Labianca

2003 "Group social capital and group effectiveness: The role of informal socializing ties." Academy of Management Journal, (forthcoming)..

O'Reilly, C. A., III, D. F. Caldwell, and W. P. Barnett

1989 "Work group demography, social integration, and turnover." Administrative Science Quarterly, 34: 21-37.

Pelled, L. H., K. M. Eisenhardt, and K. R. Xin

1999 "Exploring the black box: An analysis of work group diversity, conflict, and productivity." Administrative Science Quarterly, 44: 1-28.

Pfeffer, J.

1982 "Organizational demography." In L. L. Cummings and B. M. Staw (eds.), Research in Organizational Behavior, 5: 299-357. Greenwich, CT: JAI Press.

1985 "Organizational demography: Implications for management." California Management Review (1), 28: 67-81.

Polzer, J. T., L. P. Milton, and W. B. Swann, Jr.

2002 "Capitalizing on diversity: Congruence in small work groups." Administrative Science Quarterly, 47: 296-324.

Portes, A., and J. Sensenbrenner.

1993 "Embeddedness and immigration: Notes on the social determinants of economic action." American Journal of Sociology, 98: 1320-1350.

Reagans, R. E., and B. McEvily

2003 "Network structure and knowledge transfer: The effects of cohesion and range." Administrative Science Quarterly, 48: 240-267.

Reagans, R. E., and E. W. Zuckerman

2001 "Networks, diversity and performance: The social capital of R&D teams." Organization Science, 12: 502-518.

Reagans, R. E., E. W. Zuckerman, and B. McEvily

2004 "Two holes in one? Information and control in the analysis of structural advantage." Working paper, Columbia Business School, Columbia University.

Sørensen, J. B.

1999 "The use and misuse of the coefficient of variation in organizational demography research." Working paper, Graduate School of Business, University of Chicago.

2000 "Organizational heterogeneity and job turnover." American Sociological Review, 65: 298-310.

2003 "The organizational demography of racial employment segregation." Working paper, MIT Sloan School of Management.

Sparrowe, R., R. C. Linden, S. J. Wayne, and M. Kraimer

2001 "Social networks and the performance of individuals and groups." Academy of Management Journal, 44: 316-325.

Tajfel, H.

1981 Human Groups and Social Categories: Studies in Social Psychology. Cambridge, UK: Cambridge University Press.

Tajfel, H., and J. Turner

1986 "The social identity of intergroup behavior." In S. Worchel and W. G. Austin (eds.), Psychology and Intergroup Relations: 7-24. Chicago: Nelson-Hall.

Tsui, A. S., and B. A. Gutek

1999 Demographic Differences in Organizations: Current Research and Future Directions. Lexington, MA: Lexington Books.

Turner, J.

1987 Rediscovering the Social Group: A Social Categorization Theory. Oxford, UK: Basil Blackwell.

Wasserman, S., and K. Faust

1994 Social Network Analysis: Methods and Applications. Cambridge, UK: Cambridge

University Press.

Westphal, J. D., and L. P. Milton

2000. "How experience and network ties affect the influence of demographic minorities on corporate boards." Administrative Science Quarterly, 45: 366-398.

White, H.

2002 "Strategies and identities by mobilization context." Working paper, Department of Sociology, Columbia University.

Williams, K. Y., and C. A. O'Reilly, III

1998 "Demography and diversity in organizations: A review of 40 years of research." In B. M. Staw and L. L. Cummings (eds.), Research in Organizational Behavior, 20: 77-140. Greenwich, CT: JAI Press.

Zenger, T. R., and B. S. Lawrence

1989 "Organizational demography: The differential effects of age and tenure distributions on technical communication." Academy of Management Journal, 32: 353-376.

Zuckerman, E. W., R. E. Reagans, and B. McEvily

2004 “On firmer ground: Identifying the causal impact of social structure on performance in a contract R&D firm.” Working paper, Tepper School of Business, Carnegie Mellon University.

• Thanks to Joel Brockner, Raymond Horton, Jesper Sørensen, Toby Stuart, and participants in workshops at the University of Chicago Graduate School of Business, Goizueta Business School, the Stern School of Business and Wharton for their comments and suggestions. We take full responsibility for remaining errors.

Table 1

Gender, Race, Education, and Tenure by Function*

	Gender		Race		Education				Tenure			
Function	Male	Female	White	Nonwhite	H.S.	B.A.	M.A.	Ph. D.	<5 Years	5-10 Years	> 10 Years	Total N
Analytical	13	3	13	3	3	7	1	3	**5**	6	**5**	16
Services	(-.1)	(.3)	(-.4)	(1.5)	(-.5)	(.2)	(-.6)	(1.2)	**(-1.6)**	(1.1)	**(2.0)**	
Applied	20	2	19	3	**1**	9	4	**7**	14	6	2	22
Science	(.4)	(-.8)	(-.3)	(.9)	**(-2.1)**	(-.3)	(.8)	**(3.1)**	(.1)	(.3)	(-.5)	
Business	8	**7**	14	1	6	5	2	0	9	5	1	15
Services	(-1.3)	**(3.0)**	(.1)	(-.2)	(1.1)	(1.0)	(1.1)	(-1.2)	(-.1)	(.7)	(-.7)	
Delivery	8	0	8	0	0	5	2	0	5	1	2	8
	(.5)	(-1.1)	(.2)	(-.8)	(-1.4)	(1.0)	(1.1)	(-.9)	(.0)	(-.7)	(.9)	
Engineering	6	1	7	0	2	4	1	0	6	0	1	7
	(.0)	(-.1)	(.2)	(-.7)	(.0)	(.4)	(.1)	(-.9)	(.8)	(-1.3)	(.1)	
Materials	23	2	24	1	**14**	7	2	**0**	19	5	**0**	25
	(.4)	(-1.0)	(.2)	(-.7)	**(2.8)**	(-1.1)	(-.6)	**(-1.6)**	(1.0)	(-.3)	**(-1.8)**	
Product Life	6	1	7	0	1	6	0	0	4	1	2	7
Prediction	(.0)	(-.1)	(.2)	(-.7)	(-.7)	(1.5)	(-1.0)	(-.9)	(-.2)	(-.5)	(1.1)	
Total N	84	16	92	8	27	43	12	10	62	24	13	100

*Values in cells are the frequencies in each category. Standardized residuals are in parentheses. Bolded characters indicate a standardized residual that is significant at the $p < .10$ level.

Table 2

Analyses of Attribute Consolidation[*]

	Discriminant Analysis		Logistic Regression		
	(I) Function	(II) Education	(III) Female	(IV) Nonwhite	(V) Tenure
Dependent Variable					
Constant				-2.80	5.49
				(2.11)	(.522)
Female	F = 2.945	F = 4.441	___	1.16	-.128
				(.972)	(1.05)
Nonwhite	F = 4.421		-1.15	___	-.025
			(.912)		(1.34)
High School			**-2.00**	1.26	-.483
			(.674)	(1.24)	(.911)
Bachelor			___	___	___
Masters			.114	-.854	-1.28
			(1.20)	(.955)	(1.16)
Doctorate	F = 3.202		-.203	-.222	-.777
			(1.19)	(1.19)	(1.24)
Tenure	F = 2.391		-.014	.008	___
			(.087)	(.112)	
Model fit	Wilks' Lambda = .463	Wilks' Lambda = .869	Cox & Snell $R^2 = .114$	Cox & Snell $R^2 = .031$	$R^2 = .016$

[*]Values in columns III-V are unstandardized coefficients. Standard errors are in parentheses. Bolded values indicate a coefficient that is significant at the $p < .01$ level

Table 3

Descriptive Statistics and Correlation Matrix											
Variable	Mean	S.D.	1	2	3	4	5	6	7	8	9
1. Number of previous completions	.731	1.349									
2. Cumulative labor hours devoted to project*	3.744	1.872	.630								
3. Members' shared prior-year experience	.514	.080	.017	.036							
4. Members' shared concurrent experience	.032	.056	.376	.505	.213						
5. Team size*	1.557	.762	.535	.803	-.069	.282					
6. Function-based diversity	.493	.178	.295	.300	-.267	-.048	.293				
7. Tenure-based diversity	4.981	2.265	-.058	-.133	-.026	-.069	-.274	.159			
8. Mean tenure	5.941	3.295	-.181	-.272	-.045	-.209	-.115	-.097	-.377		
9. Internal density*	.186	.134	-.171	-.200	.380	-.050	-.126	-.392	-.164	.150	
10. External range*	.272	.030	.215	.313	.216	.148	.361	.034	-.151	.007	.383

* logged values.

Table 4

Fixed Effects Regression Models of Team Social Capital on Demography and Controls*

Predictor	(I) Logged internal density	(II) Logged external range
Constant	-.164•	.225•
	(.011)	(.002)
Team members' shared prior-year experience	.840•	.058•
	(.016)	(.003)
Team members' shared concurrent experience	.016	.027•
	(.023)	(.004)
Team size/100	.388•	.111•
	(.060)	(.012)
(Team size/100)^2	-.012•	-.0009•
	(.001)	(.0003)
Diversity		
Function-based	-.143•	.011•
	(.008)	(.001)
Tenure-based/100	-.237•	.054•
	(.061)	(.012)
Mean tenure/100	-.224•	.001
	(.027)	(.005)
Model fit		
N (team-days)	10,554	10,554
R-squared	.817	.849
Adj. R-squared	.793	.829

• $p < .05$.

* Standard errors are in parentheses

Table 5

Log-Logistic Continuous Time Hazard Models of Project Duration*

Predictor	(I) Individual dummies	(II) Controls	(III) Internal density	(IV) External Range	(V) Network structure
Constant	2.953•	1.938•	1.753•	2.934•	2.574•
	(.091)	(.341)	(.346)	(.400)	(.417)
Number of previous Completions		.103•	.102•	.103•	.102•
		(.029)	(.028)	(.028)	(.027)
Cumulative labor-hours devoted to project		.010	.012	.021	.020
		(.039)	(.039)	(.039)	(.039)
Team members' shared prior-year experience		.928•	1.847•	1.264•	1.978•
		(.407)	(.528)	(.425)	(.529)
Team members' shared concurrent experience		.429	.394	.388	.362
		(.710)	(.695)	(.694)	(.684)
Logged team size		.450•	.477•	.444•	.467•
		(.170)	(.170)	(.168)	(.168)
Diversity					
Function-based		.327	.117	.258	.090
		(.261)	(.265)	(.261)	(.265)
Tenure-based		.014	.010	.017	.013
		(.018)	(.018)	(.018)	(.018)
Mean tenure		.006	.004	.007	.005
		(.009)	(.009)	(.009)	(.009)
Network structure					
Logged internal density			-1.052•		-.889•
			(.347)		(.348)
Logged external range				-4.572•	-3.625•
				(1.285)	(1.255)
Model fit					
N of projects	1,518	1,518	1,518	1,518	1,518
N of completions	785	785	785	785	785
N of project days	10,554	10,554	10,554	10,554	10,554
Log likelihood	-980.32	-957.98	-951.62	-953.34	-948.85

• $p < .05$

* Robust standard errors are in parentheses.

Figure 1. Causal structure linking demographic diversity to network variables and team performance.

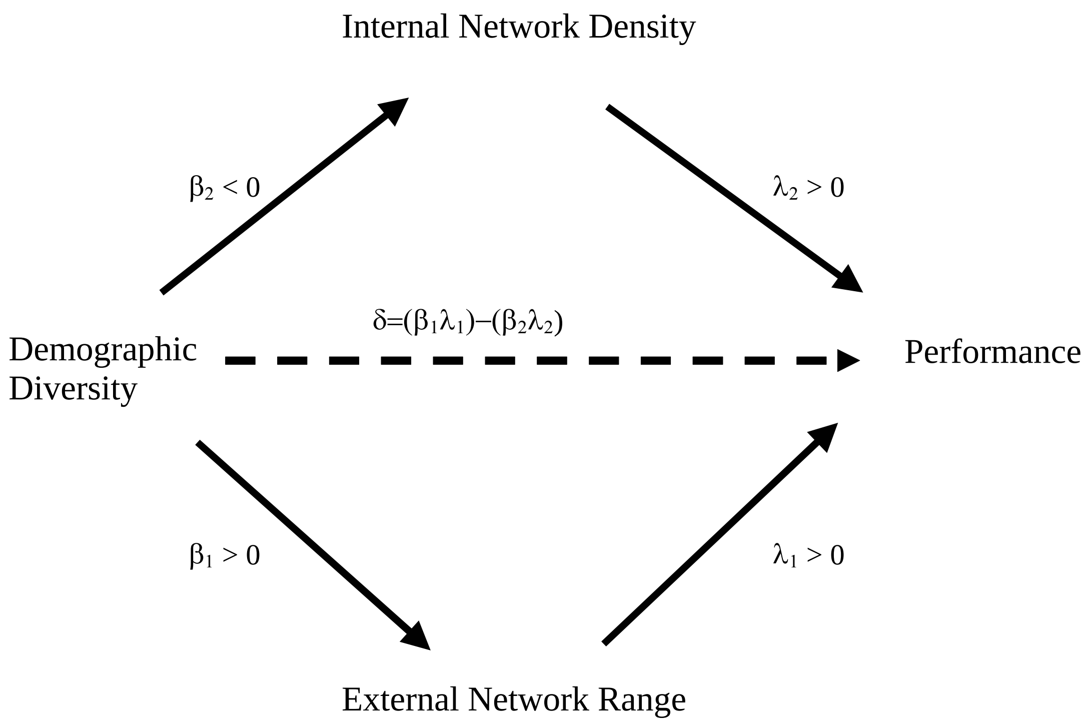

Figure 2. Local versus global structural holes.

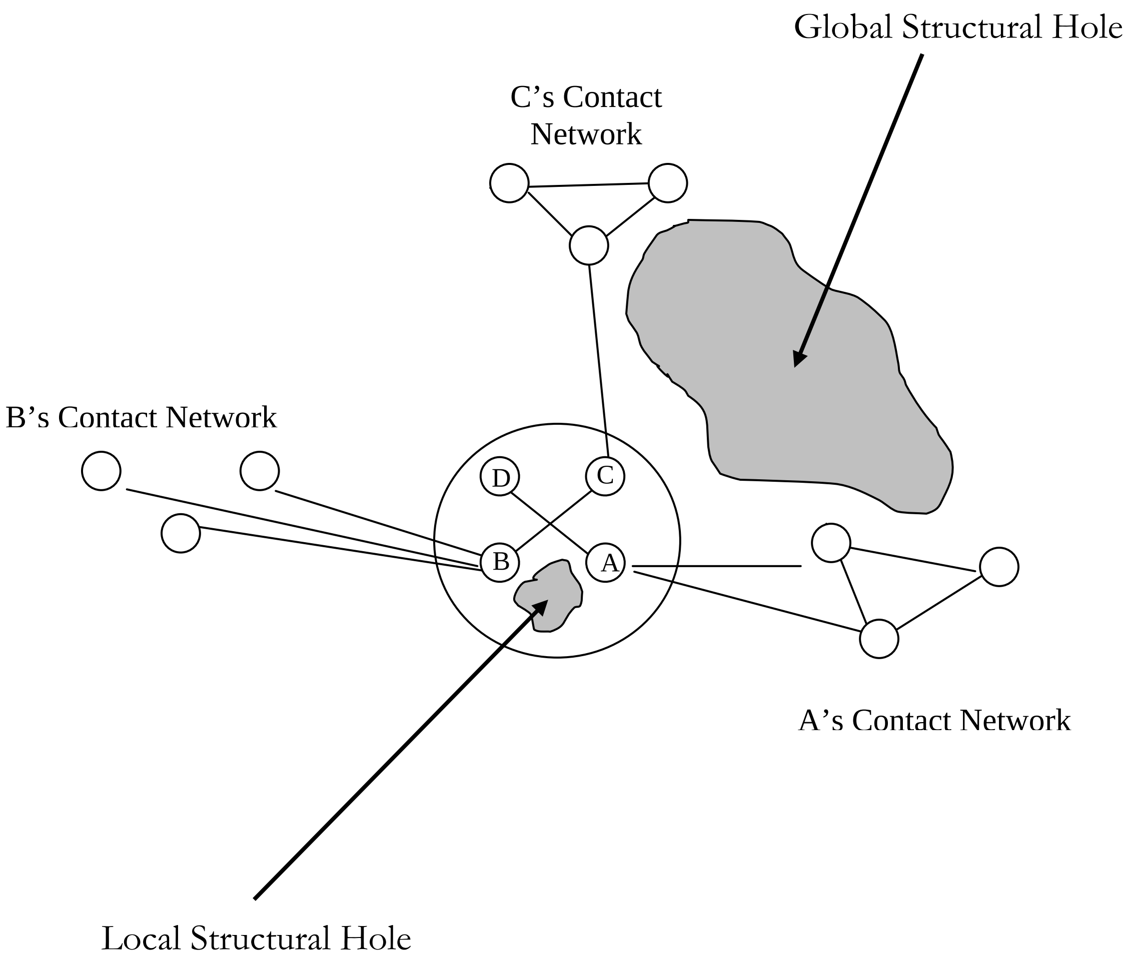

Figure 3. Estimated survival function of projects for varying levels of internal density and external range.*

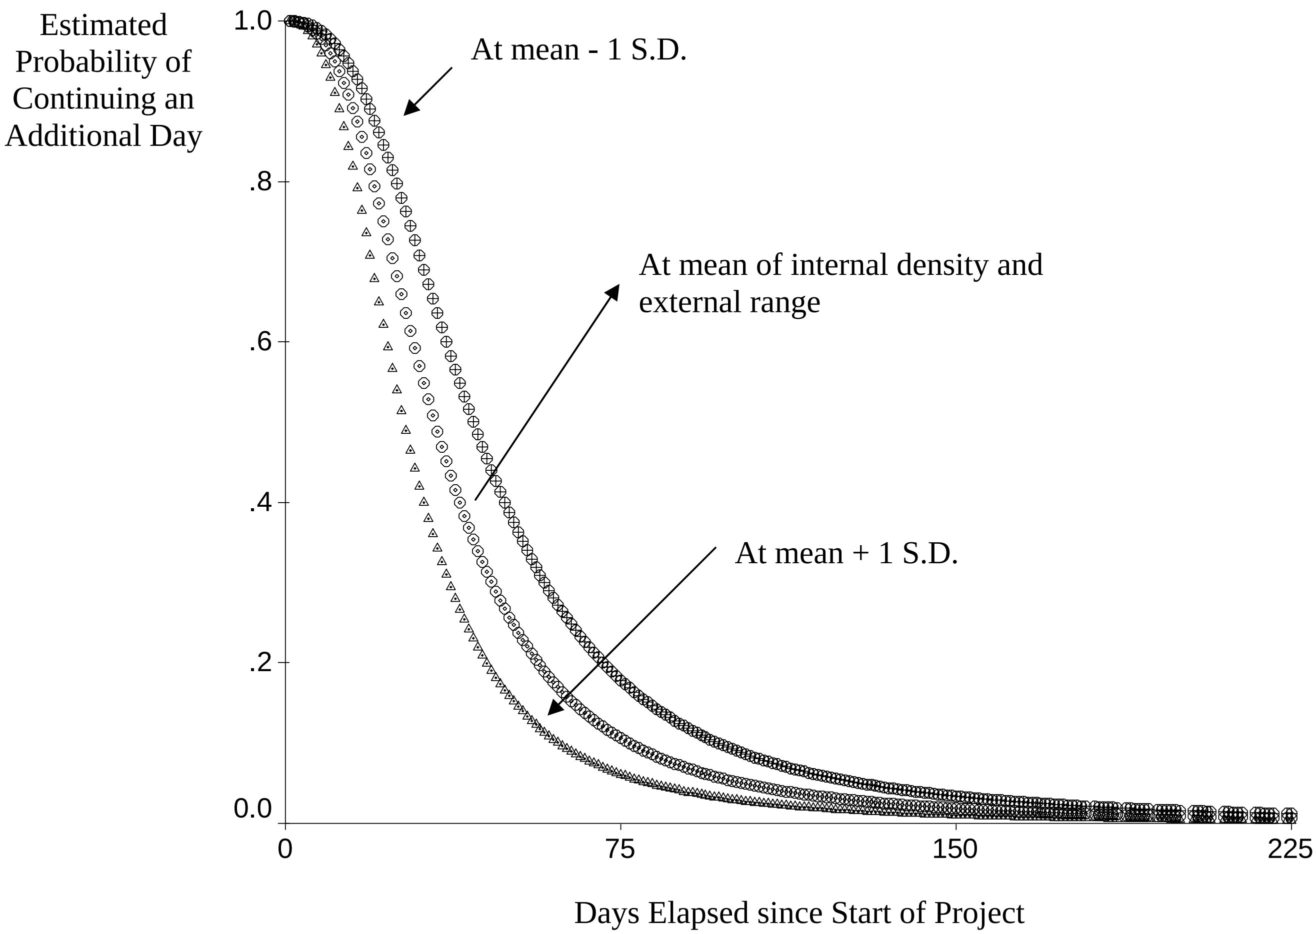

*Based on table 5, model V.

Figure 4. Estimated number of additional project completions and labor hours saved by increasing internal density and external range by 1 S.D. from mean*

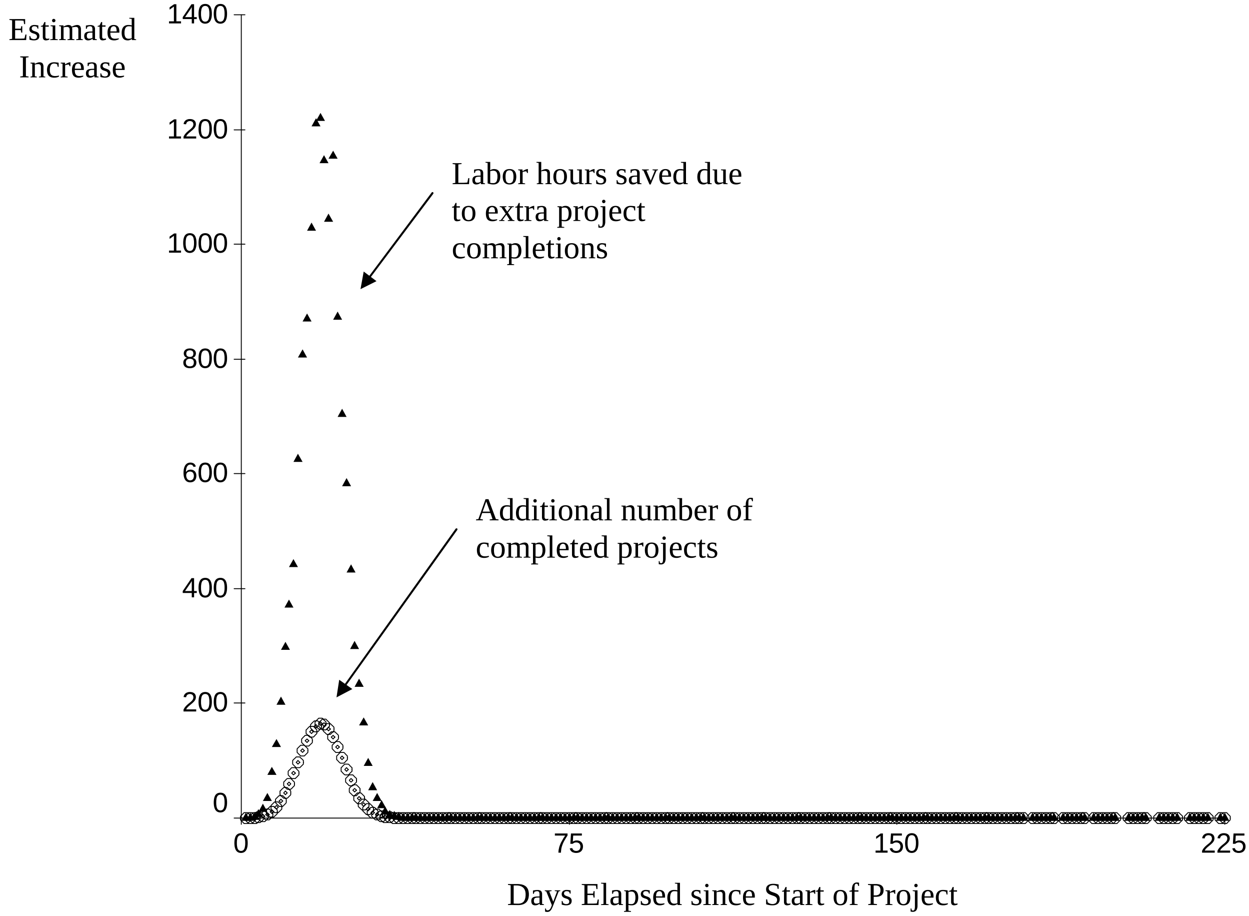

*Based on model V in table 5 and figure 3. Each circle indicates the additional number of the total 1,518 Malibu projects that are estimated to end on that date if internal density and external range are raised from the mean to one standard deviation above the mean. Each triangle reflects the number of estimated additional project completions multiplied by the mean number of labor hours that are devoted to projects of the given number of days elapsed since the start of the project.