PROGRESSING TO THE CENTER:
THE ANTECEDENTS AND CONSEQUENCES OF LATERAL AUTHORITY

Linus Dahlander
Imperial College London
SCANCOR Fellow
Stanford University
531 CERAS Building
Stanford, CA 94305-3084
linusd@stanford.edu

Siobhan O'Mahony
Assistant Professor
University of California, Davis
Graduate School of Management
One Shields Avenue
Davis, CA 95616
somahony@ucdavis.edu

October 28, 2008

We thank Beth Bechky, Paola Criscuolo, Heather Donnelly, Karim Lakhani, Fiona Murray, Jim March, Woody Powell, Ammon Salter, and the participants of the Stanford Distributed Innovation Conference and the HBS User and Open Innovation Conference for their helpful comments and insight.

PROGRESSING TO THE CENTER:
THE ANTECEDENTS AND CONSEQUENCES OF LATERAL AUTHORITY

ABSTRACT

Post-bureaucratic forms of organizing are theorized to rely upon lateral as opposed to vertical authority, but few have studied how lateral authority operates in practice. With a longitudinal, multi-network study of a mature open source project, we predict what leads individuals to gain lateral authority over collective work. While technical contributions are initially important, coordination work become more critical at a subsequent stage. After gaining authority, individuals significantly increase the effort spent coordinating project work. By specifying the antecedents and consequences of lateral authority, our research refines our theoretical conception of how knowledge work in project and communities is coordinated.

Word Count = 100

Theorists have long predicted that project and networked based forms will rely less upon traditional lines of vertical authority and more upon lateral modes of authority to achieve collective work outcomes (Barley and Kunda, 2001; Daft and Lewin, 1993; Miles and Snow, 1986; Powell, 1990; Pinchot and Pinchot, 1993; Romanelli, 1991). Despite tremendous interest in the topic, little research has examined the antecedents of lateral authority, and its consequences for the coordination of knowledge work. This research examines how lateral authority is created in the context of project work; the behaviours that enable individuals' to progress to successive positions of lateral authority; and how progression, once achieved, affects individuals' future involvement in the project.

We argue that progression can be achieved in project-based organizations that reward people with greater authority over collective work even though they do not gain authority over other individuals. In this manner, individuals can "progress toward the center" as opposed to "up a career ladder". Indeed, while this notion of lateral authority may be a critical parameter that distinguishes project work in communities or networks from project work occurring in organizational hierarchies (Powell, 1990; Simon, 1976), our current theoretical conceptions of lateral authority are incomplete. As a result, our ability to understand new forms of post-bureaucratic organizing work has been limited (Barley and Kunda, 2001).

Projects that depend upon contributions from many types of experts and specialists are less likely to take place in a single organization (Powell et al, 1996; von Hippel, 2005; von Hippel and von Krogh, 2003), but still require individuals at the center of the project to coordinate the actions of contributors who may not share a common organization (Hargadon and Sutton, 1997; Obstfeld, 2005; Long-Lingo and O'Mahony, 2008). In the worlds of science, technology and innovation, individuals frequently contribute to projects that span organizational boundaries. For example, individuals participating in technical standards committees or open source projects often do so to obtain both individual and firm benefits (Rosenkopf et al, 2001; Fleming and Waguespack, 2007; O'Mahony and Ferraro, 2007).

However, the opportunities available to those who take part in such efforts are likely to differ depending upon an individual's position in the project's network (Jones, Hesterly, Borgatti, 1997; Fleming and Waguespack, 2007). The question is – why do some individuals progress to positions of greater authority and not others? To answer this question, we analyzed the antecedents of lateral authority: the behaviours that were most likely to predict progression. But does gaining more authority actually affect individual behavior? To understand the *consequences* of progression, we analyzed how individual behavior changed once individuals gained lateral authority. We argue that understanding how authority is enacted is critical because it affects the nature of individual's subsequent contributions to the project. However, addressing both questions requires an alternative notion of progression.

Traditionally, when individuals progress in organizations that employ internal labor markets, they move through positions that grant them increasing degrees of responsibility (Althauser, 1989; Doeringer and Piore, 1971; Osterman, 1984). In bureaucratic organizations, this often includes vertical authority over individuals in the same system (Marcson, 1960). More recently, career theorists have argued that traditional notions of progression make several assumptions that no longer reflect workplace realities: (1) stable organizations in stable environments, (2) an intra-organizational focus, (3) career mobility from the standpoint of hierarchical positions, and (4) mobility within large organizations (Arthur, 1994). Thus, career theorists have called for new theories of progression – arguing that competence, identity, achievement, and community may become more critical to a conception of progression that is not dependent upon traditional notions of hierarchy (Arthur, 1994; Arthur and Rousseau, 1996; Zabusky, 1997; Ibarra, 2004). To date, few theorists have explored empirically how progression might occur within a project, community, or network that transcends organizational boundaries (with the recent exception of Lam, 2005 and O'Mahony and Bechky, 2006).

To understand what affects individual progression in projects nested in communities, we examined the importance of technical problem solving, collaboration networks and coordination behaviours to progression outcomes. The communities of practice and occupational community literatures suggest a range of behaviors such as sharing knowledge and answering questions that should help others (van Maanen and Barley, 2004; Brown and Duguid, 1991; 2001; Lave and Wenger, 1991; Wenger, 1998, 2000) and thus predict progression. Within technical communities, the ability of an individual to solve technical problems and make quality contributions should also be important to progression (von Hippel, 2005; O'Mahony and Ferraro, 2007; Fleming and Waguespack, 2007). Yet, project work that is disaggregated into modular subsections only to become integrated into a coherent whole at a later stage is also likely to require coordination work (Lawrence and Lorsch, 1967; Thompson, 1967; Gittell, 2000; Baldwin and Clark, 2006; MacCormack, Rusnak, Baldwin, 2006) which may also be associated with progression.

With a longitudinal (1999 – 2005), multi-network study of one mature community managed open source software project, we tease apart these explanations by examining the behaviors that predict which individuals acquired lateral authority at two successive stages of progression: (1) becoming a member of the foundation and (2) being elected to the board of directors. A community managed open source project is a distributed group of individuals developing open source code with the support of the Internet that is independent from employment relations, pluralistic, and fostered by decentralized decision making and autonomous representation (O'Mahony, 2007). Our longitudinal data set offers a unique opportunity to view cohorts moving through two successive stages of a pipeline. This turns out to be important as we find that different mechanisms are responsible for different stages of progression and that individuals respond differently to each stage. While technical problem solving predicts an individual's progression from project contributor to "member" at the first stage, it does not explain progression to becoming a board of director. Centrality in communication networks and involvement in coordination work is critical to become elected.

After analyzing why some individuals gain lateral authority and others do not, we explored the consequences of gaining authority by examining how individuals responded to their new positions. Contrary to our informants' expressed belief that progression would not change their behavior, we found an immediate post-progression effect: project members contributed significantly more effort to the project after gaining lateral authority. At both levels of progression, individuals redoubled their efforts –most notably with respect to coordination work.

This research makes three theoretical contributions to our understanding of how knowledge work is coordinated in community based projects. First, we refine the distinction between two different types of authority - authority over collective work (lateral authority) versus authority over individuals (vertical authority). Our conception of lateral authority helps explain how progression without formal hierarchy can occur in project forms of organizing – by progressing to the *center*. Second, a two stage empirical model allows us to show how individuals progress through such social systems over time and to parse how progression mechanisms differ by stage. Activities that foster coordination on the project become more important as individuals gain authority over core project activities. Third, we show that lateral authority may be an under appreciated mechanism that fosters not only coordination work, but also individual contributions to a project – which in turn improves the chances of such forms surviving. Taken together, by explicating the mechanisms that facilitate the antecedents and consequences of lateral authority for both individuals and projects, we further our understanding of how knowledge work is coordinated in community forms.

EXPERTISE AND PROGRESSION

Progression in traditional organizations usually introduces greater responsibility and challenge through progressive sequences in a career ladder. At higher levels, progression in bureaucratic organizations often includes managerial authority over other individuals. However, career theorists have long recognized that scientific or technical experts may seek progression without assuming managerial responsibilities. Beginning with Gouldner's distinction between 'local' and 'cosmopolitan' orientations (1968), theorists have recognized that individuals with specialized expertise often demonstrate an academic or scientific orientation to an external community more so than they do to their proximate or 'local' community (Marcson, 1960; Ritti, 1968; van Maanen and Barley, 1984). Scientific and technical professionals therefore often seek progression in terms of knowledge, acumen and skill rather than advancement into managerial positions (Zabusky and Barley, 1996). For if technical experts are rewarded by managerial promotion, they risk becoming separated from their chosen discipline or knowledge base (Golder and Ritti, 1967) and can face obsolescence of their chosen skill (Rothman and Perrucci, 1970). Academics and practitioners alike have wrestled with the managerial challenge of how to reward experts who do not seek managerial authority but instead seek greater autonomy (Bailyn, 1991).

One solution to this challenge is a “dual” career track – which allows technical experts to achieve greater autonomy, prestige and rewards without assuming managerial responsibilities (Bailyn, 1991; Katz, Tushman and Allen, 1995). Individuals who sought managerial advancement could progress through the managerial track, while technical experts who wanted to advance in their respective fields could do so without managing people (Katz et al, 1995). Dual career ladders could thus provide equivalent status and rewards in two distinct, but parallel, hierarchies (Allen and Katz, 1986). “Instead of greater authority, [technical experts] are rewarded with greater freedom to engage in their specialities” (Kornhauser, 1962: 205). However, the degree to which dual ladder career systems actually provided equivalent rewards was much debated (Goldner and Ritti, 1967; Ritti, 1971). Indeed, firms often diluted the ‘separate but equal’ technical ladder by using it as an opportunity to provide rewards to individuals deemed unsuitable for the managerial track (Allen and Katz 1986).

This literature was inspired by the “challenge of managing scientific and technical expertise” in industrial research labs in the 1960s and 70s. With the rise of industrial research and development, came the recognition that individuals pursuing highly specialized work might require a different reward system and career path. These scholars may well have been at the forefront of understanding a key challenge for the 21st century (Powell and Snellman, 2004) - how to manage experts in knowledge intensive industries without depending upon managerial authority. However, this literature’s contribution to understanding the organization of knowledge work has not been thoroughly mined. Because scholars in this tradition focused narrowly on the management of research and development in industrial labs, the “challenge” of managing experts was viewed as a dyadic motivation/reward problem *between* managers and experts. How work was coordinated *among* experts was not explored. Managers and theorists understood that scientific and technical workers sought autonomy and recognition in their respective communities but how this would affect work among likeminded peers was not examined. What was left under explored was the role of lateral authority.

Barley and colleagues took a step in this direction by arguing that the infusion of technical work in a broad class of occupations would call for a horizontal as opposed to a vertical organization of work (Barley, 1995; Whalley and Barley, 1997; Zabusky and Barley, 1996; Zabusky, 1997; Barley and Kunda, 2001; Barley and Kunda, 2004). "Horizontally organized work…is focused less on power and more on expertise. In a horizontal system, different groups of practitioners jointly contribute their distinctive efforts to the execution of work tasks...collaboration rather than command, is the key to getting work done." (Zabusky, 1997: 130). This conception assumed that somehow, experts would need to figure out a way to collaborate *with each other* without the authority relations implied by a vertical hierarchy. Yet, the antecedents, mechanisms and consequences associated with lateral authority remained unexplored.

Our definition of lateral authority allows individuals to gain increased control and responsibility over *collective* work without necessarily gaining authority over other individuals. This conception contrasts with the literature on technical career ladders, which merely promised scientists and engineers greater autonomy over their *own* work. For example, contributors to technical committees and projects can accrue greater influence, control and authority over the project, but not over the people contributing to the project (Fleming and Waguespack, 2006; O'Mahony and Ferraro, 2007). In this manner, people can acquire greater decision rights over coordinating a project's collective work without necessarily assuming responsibility for managing subordinates. However, the implications of this notion of progression for coordinating knowledge work have not been explicated.

As opposed to examining 'separate but equal' hierarchies, we explore how progression occurs in a project that depends solely upon lateral authority where progression is best conceptualized as movement toward the center. Figure 1 contrasts 'progression to the center' with a traditional pyramid, where a smaller number of people at the top are accountable for managing a larger number of people beneath them.

- Insert Figure 1 about here -

We argue that understanding how progression operates in lateral as opposed to vertical systems of authority will enable more generalizable theory building as to how knowledge work in community and project forms is coordinated. This type of progression is likely to apply to other types of experts and specialists whose contributions are vital to the creation of new knowledge or innovations (e.g. Obstfeld, 2005; Hargadon and Bechky, 2006). Our hope is that by examining the antecedents and consequences of lateral authority in a large community managed open source project, we can offer a more specific and generalizable understanding of the relationship between lateral authority and the coordination of knowledge work. Because open source projects exist independent of any one organization, we draw upon literatures that examine knowledge work in communities to identify factors that could explain progression in community forms.

KNOWLEDGE WORK IN COMMUNITIES

The literature on occupational communities and communities of practice examines how individuals solve problems and share knowledge within informally constituted networks of shared practice, that are not limited by organizational boundaries (van Maanen and Barley, 1984; Brown and Duguid, 1991; 2001; Wenger, 2001; Bechky, 2003a, b). Thus, this literature is appropriate to understanding how individual contributions affect progression on an open source community project. In drawing upon this literature, we identified three primary behaviors that could predict progression within a community project: the degree to which individuals engaged in: (1) technical problem solving; (2) collaboration networks; or (3) coordination work.

Technical Problem Solving. Individuals engage in communities of practice in order to learn from each other (Wenger, 1998) and help close the gap between knowing and doing (Brown and Duguid, 1991). If communities depend upon the technical contributions of their members in order to continue to attract participants and survive (Butler, 2001), then individuals who contribute more technical contributions should be more likely to progress to positions of greater responsibility (von Hippel and von Krogh, 2003). Consistent with this view, theorists tend to assume that open source communities operate in a meritocratic manner (Lee and Cole, 2003; Kogut and Metiu, 2001). Merit in this sense, refers to technical contributions of software source code – where one's real technical prowess is both discernible and usable by others (Raymond, 1999).

There are several reasons why an individual's technical contributions, once accepted by the community are likely to improve one's chance for progression. First, other community participants are appreciative of individuals who devote time working in the community. For example, contributions that are included in the code base of an open source software project, have passed an intense peer review process very similar to that of an academic paper. Thus, individuals with demonstrated effort to a project may be more likely to gain the trust of other project members. Second, individuals may need to demonstrate a track record of technical knowledge in order to be granted the authority to make decisions that will affect the entire project. Project members may perceive individuals who make more technical contributions as more knowledgeable about project dependencies and thus a more credible decision-maker on project affairs.

> H1: The more technical contributions an individual has accepted by a project's peers, the greater the probability of progressing to positions of increasing lateral authority.

On the other hand, offering a new technical contribution to a project may be the result of an individual satisfying their own needs and desires (Lakhani and von Hippel, 2003), rather than demonstrating the ability to resolve problems that are of long standing interest to other project members. While creating new content can be a creative and enjoyable task, fixing and maintaining what has already been developed is considerably less so. Furthermore, solving existing project problems may require more extensive knowledge of project dependencies and thus more complex problem solving skills (Baldwin and Clark, 2006; MacCormack, Rusnak, and Baldwin, 2006).

Building on this reasoning, a number of scholars have suggested that individuals accumulate power to make critical decisions on a project based on their ability to solve complex technical problems (von Krogh, Spaeth and Lakhani, 2003). By demonstrating their ability to resolve problems that could be detrimental to the project as a whole as opposed to new contributions that may only affect a few, project contributors may acquire more status (e.g. Stewart, 2005). Thus, we suggest that the more pre-existing technical problems an individuals can resolve, the more likely he or she will be to progress to positions of increasing authority.

> H2: The greater an individual's involvement in solving pre-existing technical problems, the greater the probability of progressing to positions of increasing lateral authority.

Not everyone may benefit directly from new contributions or even notice the resolution of a technical problem, but as many scholars have argued, one of the critical benefits that communities offer is the ability to learn from others (van Maanen and Barley, 1984; Brown and Duguid, 1991; 2001). By engaging in on-going collective sensemaking and problem solving, community members learn from *modus operandi* (the way a task appears in progress) as opposed to deconstructing *opus operatum* (what is visible after the task is finished) how a task should be performed (Bourdieu, 1973; Brown and Duguid, 1991). Thus, posing and answering questions as the work is performed is essential to fostering learning (Bechky, 2003a,b). In online communities, engagement in conversational threads is the primary way in which individuals learn from each other and share knowledge (Yates et al, 2003; Orlikowski and Yates, 1994). In this setting, a conversational thread is "a stream of conversation in which successive contributions continue a topic, following an initial contribution which introduces a new topic" (McDaniel, Olson and Magee, 1996: 41).

Because a discussion thread allows others to build upon it, those who initiate such threads can shape the community's conversational agenda and potentially gain the respect of their peers. Three recent studies of the emergence of informal leaders in online communities found that 'leaders' posted more messages than non-leaders (Cassell et al, 2005; Misiolek and Heckman, 2005; Yoo and Alavi, 2004). Thus, we suspect that those who initiate technical discussions may become more likely to progress to positions of greater authority. At the same time, we recognize that initiating a discussion can also reveal one's ignorance to others. Thus, those who respond to queries may also be rewarded with more authority over the project.

> H3a: Individuals that initiate technical discussions will have a greater probability of progressing to positions of increasing lateral authority.

> H3b: Individuals that respond to technical discussions will have a greater probability of progressing to positions of increasing lateral authority.

Collaboration Networks. Even though project work may be geographically distributed (Hinds and Kiesler, 2003), individuals may still collaborate at smaller sub-project levels to make contributions to the project (e.g. Baldwin and Clark, 2006). As is true in other contexts, the structure of collaboration networks may affect which individuals progress to positions of authority. We suspect that individuals who maintain more sub-project collaborations may develop a broader understanding of the project, its constraints and its resources. Engaging in technical collaborations, rather than just discussing technical issues, may be a more concrete and visible way in which individuals provide value to the project.

> H4: Individuals that collaborate on a larger number of sub-projects will be more likely to progress to positions of increasing lateral authority.

In forming judgements and beliefs about others, individuals often rely on the information they receive through their social networks (Podolny, 1993) especially when firsthand information is not available – which may be the case when project members are distributed. It is generally accepted that individuals who hold more structurally central positions within a social network are more likely to enjoy greater influence than those on the periphery of the network (Brass, 1984; Gould 2002). In organizations, individuals with more favorable structural network positions are more likely to become promoted (Burt, 1992; Katz et al., 1995; Podolny and Baron, 1997).

However, in organizations, career mobility decisions are made by those in higher positions of vertical authority. What happens when promotion decisions are made by one's peers? It is possible that the same effect holds, but different mechanisms may be at play. Those that are structurally central in a project's communication networks may be better able to achieve visibility among their colleagues and have the opportunity to demonstrate their dedication and effort to the community. In contrast, individuals on the fringe of the network may find it more difficult to gain the attention of other project contributors and earn their respect.

> H5: Individuals who are more structurally central in a project's communication network will have a greater probability of progressing to positions of increasing lateral authority.

Coordination Work. Coordination has been defined as the additional effort individuals expend when attempting to achieve collective goals over and above the effort they would expend to achieve that goal individually (Malone and Crowston, 1994) or as the act of managing interdependencies to achieve collective goals (Malone and Crowston, 1990; Crowston, 1997). Coordinating individual contributions to achieve a collective outcome is an ongoing challenge for all organizations (March and Simon, 1958; Lawrence and Lorsch, 1967; Ouchi, 1980; Bechky and Okuhysen, 2008) but can be an even greater challenge when project members are distributed and rarely meet.

Large technical projects are often disaggregated into more manageable modules or sub-projects (Baldwin and Clark, 2006). Yet, the creation of many subprojects can fragment knowledge if efforts to coordinate individual efforts are not made. The trick is to integrate sub-project modules into a coherent whole once an individual's work is accomplished (Mintzberg, 1979). To foster integration work, technical tools can prevent individuals from working on the same thing at the same time and enable the tracking of project changes and activities. Such systems foster a shared understanding of the work that has been *completed* and provide accountability as to who accomplished what task: two elements that are critical to the coordination of work (Okhuysen and Bechky, 2008). By offering standardized interfaces to organize the receipt of contributions from different parties, these systems help manage the interdependence shared by different project subcomponents (Thompson, 1967).

However, these systems provide little visibility as to the work in progress nor do they offer information about how distant sub-projects might be connected to each other. Thus, such tools are not sufficient to foster coordination through mutual adjustment (March and Simon, 1958). When projects are dynamic and highly interdependent, they will need integration mechanisms to provide feedback across functional lines (Thompson, 1967; March and Simon, 1958). Once a project is disaggregated into sub-projects, individuals need mechanisms to maintain awareness of activities in other areas of the project that may be relevant if not interdependent with their own area. When there are few formal coordination mechanisms in place (such as on an open source project), the coordination mechanism that may become most critical is the ability to span and share knowledge across sub-community boundaries.

In organizations, boundary spanning has been found to enhance innovation (Tushman, 1977) and foster career progression (Katz and Tushman, 1983). Individuals who span subproject boundaries may help transfer knowledge; identify new ideas (Allen, 1977; Tushman, 1977; Hargadon and Sutton, 1997; Hansen, 1999), enhance community members' awareness of the broader project environment (Ancona and Caldwell, 1992), and expose unattended or unidentified project interfaces (Sosa, Eppinger and Rowles, 2007). Gittel and Weiss'(2004) study of health care service delivery suggests that boundary spanners such as healthcare case managers performed a critical coordination function that improved the quality and efficiency of care delivered in hospitals. No doubt, boundary spanning helped enable mutual adjustment to occur in an environment that depends upon rapid response. Boundary spanning may be even more critical on projects that are distributed and reciprocally interdependent. If individuals span more subproject boundaries, they may be more likely to understand the needs of the whole project and be in a position to proffer useful solutions, ideas and resources - enabling progression to positions of greater lateral authority.

> H6: Individuals that span more subproject boundaries will have a greater probability of progressing to positions of increasing lateral authority.

To help coordinate individual contributions to the project, it is important that individuals share the knowledge gained from boundary spanning with other parts of the project. In their study of the IETF, Fleming and Waguespack (2007) found that boundary spanners were more likely to assume leadership positions than brokers who were structurally unique in their position. Rather than maintaining unique information, what was important to assuming a leadership position, was that boundary spanners shared any information gleaned from their activities with others. On an open source project, this is likely to happen two ways: (1) by either initiating discussions on coordination issues or (2) by answering questions in forums devoted to the coordination of project-wide activities.

> H7a: Individuals that initiate more coordination discussions will have a greater probability of progressing to positions of increasing lateral authority.
>
> H7b: Individuals that respond to more coordination discussions will have a greater probability of progressing to positions of increasing lateral authority.

Consequences of Lateral Authority. To further theory that specifies the relationship between progression and the coordination of knowledge work, we were interested not only in the antecedents that predicted progression, but also its consequences. Moving into a position of authority may trigger the adoption of a new role that brings with it norms and expectations to be met as well as new resources with which to fulfil those expectations (Turner, 1986). Futhermore, the articulation of a clear role structure is yet another means by which coordination can occur – particularly in project forms that lack traditional hierarchical coordination structures (Bechky, 2003a).

Research on high reliability organizations shows that when there is complex, highly interdependent work, individuals use role structures to coordinate their behavior and avoid errors (Weick and Roberts, 1993; Weick, 1993). Even in temporary projects that lack stable rules and hierarchies, such as film sets, structure is provided through the continuous enactment and negotiation of role expectations (Bechky, 2003a). Progressing to a role that offers a clear basis of authority may foster an individual's ability to engage in coordination work by providing them with more information about the project; by communicating role expectations; or by simply resolving ambiguity over who has authority over what. We suspect that all of these mechanisms may prompt individuals who assume clear roles of lateral authority to take on more coordination work.

> H8**:** After progressing to positions of increasing authority, individuals will increase the effort devoted to coordination work.

RESEARCH SETTING

Barley and Kunda suggest that because projects are "a primary locus of affiliation and decision-making", they are an appropriate focus of theorizing. They argue that "a comprehensive understanding of organizing…require[s] greater attention to the logistical, technical, temporal and managerial dynamics of project life" (2001: 79). Because project forms require detailed, contextually sensitive data in order to understand them (Barley and Kunda, 2001), we first describe our research setting and explore how lateral authority and progression were operationalized in this context. Our understanding of this open source project draws from over a dozen ethnographic interviews conducted prior to conducting the quantitative analyses. We also studied numerous reports and articles on the community and examined archival information on the project's website to understand the community's history and development and guide our approach to data collection and analysis.

GNOME, the community managed open source project we study, develops a user-friendly desktop that provides a Windows-like environment for Linux, and a platform for developing applications that integrate with the rest of the desktop. We chose GNOME as our research setting for three reasons. First, the project has operated successfully for over ten years, providing us the opportunity to explore phenomena that might be less available in more nascent settings. Second, like other large open source projects, GNOME has developed a formalized governance structure, with members that are appointed, and an elected board. We can therefore follow cohorts moving through this system and the behaviours that characterize those individuals that move into positions of authority. Third, we were able to retrieve multiple, compatible longitudinal data to test our hypotheses.

Within two years of its inception, the GNOME project had attracted several thousand individuals who were interested in contributing to the project. With this rapid growth, a number of individuals began to voice concerns about the need to develop a more formalized governance structure as the ad hoc mode of governance was showing signs of strain (e.g. O'Mahony and Bechky, 2008). As one founding member recalled: "there was no formal decision making, you just had this sort of total chaos…. there was nobody that could just come in and say…let's work it out and then decide, this is what we're gonna do, because there was no way to make a decision, there was no one in charge."

After discussions on project mailing lists, the founders decided (with the help of collaborating firms) (e.g. O'Mahony and Bechky, 2008) to create a non-profit foundation that would create a governing body and help the project manage its increasing scale:

> When GNOME was a smaller project, [the founder or the community] was able to make most of the key decisions. […] We need a more structured environment to smoothly integrate new citizens into the community. The GNOME Foundation will provide this support. […] Whereas decisions in the

> past have often been made in an ad-hoc fashion and in private conversations between a small number of people, the foundation will provide a forum that is elected by the GNOME community, that is accountable to that community, and that will conduct its affairs in the open.
> http://foundation.gnome.org/about/charter/

The creation of a formal governance structure was also viewed as a way to foster coordination and decision-making within the project. As one GNOME informant recounted, "most projects have a very well-defined hierarchy, technical hierarchy. So, at some stage someone says, "Yes it goes in" or "No it doesn't"…And in GNOME that last stage is missing. Now we have the foundation."

A New System of Lateral Authority. The governance structure that project members designed created two formal successive stages of lateral authority: members and board of directors; the former a condition of receiving the latter. The question of interest is which individuals were promoted to these new roles and, once granted, how these roles were interpreted and enacted.

Members of the foundation. Becoming a member grants individuals the right to make decisions that affect the whole project – for example, the right to admit new members and decide upon software release criteria and the right to elect a board of directors. With a class of members, the project could narrow who could influence the project's future direction. It provided the community with a way to weed out contributors who would not remain committed to the project and to recognize those committed to the project long term. As one board director explained, "we get twenty people a day saying "I'm going to do this, or I want to do this, you know? People try to get involved and it's just too many people. There is no way we can deal with them. And the fact is 99% of those people disappear after two weeks….about 1% actually end up becoming valuable contributors, members of the community [that are].” At the same time, as one of the founding board directors explained, the GNOME community wanted to ensure that “we’re a really open project, we like to be open, get a lot of people involved and be sure everybody has a voice and are contributing, we do want people to be contributing.” Specifying a member class and a process for membership helped fulfil both objectives.

Board of Directors. Board of directors have the right to make decisions regarding the project’s assets and resources and the ability to represent the project to the public. Only those that are members could become a board director through a democratic election. The decision to elect their board of directors was driven by the community’s reluctance to hand over any type of authority to specific individuals. As one of the drafters of the foundation’s charter recalled, “Everyone was saying, this thing needs to be elected. Right? Folks felt like, you know, if I’m a hacker, I should have a voice in this I should be able to decide who gets to run this thing.”

The board of directors could grant companies permission to use GNOME’s brand and trademarks, push the project to close on a specific release date, fund a technical conference and shape the technical direction of the project. “It’s not really feasible to have all several hundred GNOME developers vote on every little thing that comes up, you don’t just want to have direct democracy there, so the compromise we came to was this board who would be elected by them [the members] that would sort of take care of the day to day details.” At the same time, members recognized that the board did not gain authority over individuals – a distinguishing feature of lateral authority. “It’s not like the board can tell several hundred developers what to do…we’re all just volunteers, so there’s some cat herding going on”.

We consider both forms of authority to be lateral as they provided individuals who had contributed to the project with increasing responsibility for the *project* without granting them authority over *individuals*. As the project founder emphasized, even with the creation of this new structure: "all the technical decisions are still in the hands of the developers." Contributors remained free to work on the areas that were of interest to them and to withdraw from the project at any time. Even so, the creation of any type of authority structure was viewed by many as a 'necessary evil'. As an early board of director explained, "computer programmers and particularly free software hackers do not like bureaucracy or authority or anything like that right? But at the same time, I think we're all pretty pragmatic and at some point you have to have some level of organization if you're going to get all those interests pulling in the same direction." Project contributors recognized that by creating some structure they could enhance decision-making and coordinate work among their growing members. The question, is did people's behaviour change after gaining lateral authority? What were the consequences of creating this new role for individuals and for the project?

The consequences of lateral authority. Our informants indicated that the roles of member and board of director did not change the nature of their work nor the meaning it held for them. However, informants also frequently spoke of ways in which the creation of clear-cut authority roles helped the project come to agreement on critical areas where consensus decision-making had stalled. For example, several individuals were frustrated when other project contributors continued to change the code even though it was supposed to be frozen for a major release. One individual, frustrated with the situation, asked the board of directors to shepherd the project and help coordinate a software release: "You need a Foundation, you need a group of people to say, "Hey folks, we decided it's frozen, stop fucking with it…What it took was a board of directors elected by all the hackers…saying it needs to be frozen!

…You need that much power, to be able to make these calls". It is possible that an 'authority vacuum' existed where individuals did not coordinate project work because the authority roles were unclear. Our quantitative data allows us to examine more precisely not only what predicts individual progression at two successive levels, but also how these new roles affected individuals'subsequent behavior.

METHODS

Once we had a detailed picture of how the community works, we obtained yearly information on whether individuals were granted membership or elected to the board of directors, interactions on all mailing lists, contributions of source code and bug solutions, as well as individual characteristics. We defined individuals as belonging to the community and thus, in our population pool, if they sent at least one mail to one mailing lists. Individuals were dropped from the analysis if no further interactions ensued after two years (Uzzi and Spiro, 2005). The panel used for analyses includes data from all interactions on community mailing lists and all contributions of source code between 1999 and 2004, and information about individual progression in the GNOME foundation for the period 2000 to 2005.

Dependent Variables: Lateral Authority

We examined the antecedent behaviors that affected the probability of individuals' progressing at two different levels: (1) becoming a member of the foundation and (2) being elected to the board of directors.

Progression Step 1: New Member. We measured $New\ Member_{it}$ with a binary variable taking the value 1 if an individual became a new member in the foundation, for individual *i* in year *t* by coding the available information on the website. We cross-checked website data with information from the mailing list and membership applications from 2000 to 2005 to ensure reliability. According to the project's charter, membership is granted only for "non-trivial improvement[s]" to the GNOME Project. Contributions could come in the form of "code, documentation, translations, [or] maintenance of project-wide resources"[1]. Membership must be renewed every two years, so the total number of members may change from one year to the next. Because the likelihood of renewing membership is very high once an individual becomes a member, we focus on new members only, and drop observations after an individual is subject to membership renewal. This is a conservative test of our hypotheses and the alternative approach with using renewal of membership yields similar results.

Progression Step 2: Elected Board of Director. Since only members can be considered as nominees for the annually elected board of directors, the second level of progression is conditioned on the first. As a consequence, the dataset we use to predict board directors is smaller than the one used to predict new members. To predict the probability of being elected to the Board of Directors, we measured $\textit{Elected Director}_{it}$ with a binary variable taking the value 1 if an individual became elected, for individual i in year t. We retrieved all election data from the foundation's creation in 2000 to the year 2005.

Independent variables

Our independent variables are derived from four primary sources – the community's website, mailing list archives, the Concurrent Version System (CVS) (which manages the checking in and out of software source code), and the bug database called Bugzilla (which tracks the identification and resolution of technical problems on the project).

[1] For more information, see: http://foundation.gnome.org/membership/.

To obtain a list of all individuals that participated in the community, we downloaded all the 387,626 emails sent to community mailing lists between 1999 and 2004 to a new database. This database included information about the email of the sender, date, subject, message content, who the sender responded to, conversation thread identifier, and an indication of the mailing list. Unique individuals were identified first by their email-addresses, and second by checking whether the same individual owned two mail addresses. Different ways of writing the same name were manually accounted for by checking various alternatives (compare Linus Torvalds, L. Torvalds, Torvalds, Linus). This procedure resulted in a set of 24,694 unique individuals who sent at least one mail to any of the mailing lists during the time of study.

A mailing list is a simple communication technology that broadcasts an email to a list of subscribers signed up to that particular list. There are a multitude of mailing lists targeting different areas of the project. Individuals identify the mailing lists appropriate for their queries, and avoid posting messages in multiple lists as "wide cross-posting is discouraged" (http://mail.gnome.org/mailman/listinfo). Prior research has noted the importance of mailing lists in open source communities (Kuk, 2006; Lee and Cole, 2003; von Krogh, Spaeth, and Lakhani, 2003), but has typically only investigated a fraction of all interactions in a community. Our analysis is based on the project's complete communication venues: all 137 mailing lists active between 1999 and 2004. Mailing lists that contained only automated updates with changes of source code and spam messages were deleted because they do not represent opportunities for interactions between individuals.

Because postings to a mail list are broadcasted to a list of subscribers, it is easy for newcomers to participate. When a subscriber receives a mail, she can choose to reply so that all subscribers receive the response. These data enable the tracking of all interactions within all of the community's communication networks - which constitute the predominant method for discussing and solving problems within GNOME. Initiating discussions, posting questions and responding to mailing list queries are vital contributions to the community's efforts to develop, maintain, improve and foster the usage of their software.

We also downloaded all commits of source code to GNOME's concurrent version system (CVS) to our database. CVS tracks changes in the source code. Source code that has been checked into CVS has gone through the community's peer review process where individuals evaluate the code to see if the contribution constitutes an improvement to the code base. The CVS data reveal when an individual submitted a source code change, in what directory, and the nature of the contribution. This represents a unique opportunity to track all technical contributions over time. We downloaded data for all contributions from the 1,146 individuals that contributed 1,915,085 technical contributions from 1999 to 2004. We draw from these data sources to measure three types of antecedents for progressing in the community: 1) technical problem solving, 2) collaboration networks; and 3) coordination work.

(1) *Technical Problem Solving*. We use four different independent variables to measure an individual's ability to solve technical problems in the community: commits of source code; resolution of bugs; initiation of technical discussions; and responses to technical discussions.

Number of source code commits. One measure of technical problem solving is the natural log of the *number of source code commits*$_{it}$ for individual i in year t which is calculated from the number of times an individual has checked in source code in the CVS.

Number of bugs resolved. Resolving bugs is an important technical task, which, while not as attractive as creating new code, can, if left unattended, prevent the community from releasing a final distribution of their code. We measured this activity with the natural log of the *number of resolved bugs*$_{it}$ an individual *i* had completed in year *t*.

Number of technical discussions initiated. We content coded the 137 mailing lists used by the GNOME community to distinguish different types of activities. The bulk of the mailing lists (between 96.5 – 99.5% depending upon the year) were focused on technical discussions. We measured how proactive an individual was in creating new technical discussions with the natural log of the *number of new technical discussions initiated*$_{it}$ by an individual *i* in year *t*.

Number of responses to technical discussions. To test whether individuals that responded to technical discussions were more likely to gain lateral authority, we developed a variable with the natural log of the *Number of responses to technical discussions*$_{it}$ to threads by individual *i* in year *t*.

(2) *Collaboration Networks.* We assessed the degree to which an individual collaborates with others with two measures: one measuring direct technical collaborations on source code and the other measuring an individual's centrality in the communication networks that support collaboration.

Number of technical collaborations. With CVS data, we measured an individual's *number of technical collaborations*$_{it}$ by the number of sub-projects where an individual *i* contributed source code divided by the overall number of sub-projects in year *t*. The larger the number, the greater the proportion of sub-projects to which an individual has contributed source code

– and the greater the number of collaborations. In supplementary analyses, we used degree centrality in this code collaboration network. This alternative measure correlated .85 with our measure of technical collaborations, and led to the same substantive conclusions.

Structurally central in communication network. With mailing list data, we derived information about who interacted with whom. To obtain a measure of *structural centrality in communication networks*$_{it}$ maintained by an individual *i* in year *t*, we counted the number of participants an individual had interacted with in shared threads across all of the mailing lists. Following Freeman (1979), we used degree centrality which measures the number of individuals an individual is connected to in the network. To account for the fact that the size of the community changes over the years, we normalized the measure by dividing it by the total size of the community at year *t*. Because there is a tendency for ties to weaken and dissolve over time (Burt, 2000; Soda, Usai, and Zaheer, 2004), the probability of decay increases with age. Interviews with community participants suggested that individuals that were inactive for three years or more, should be removed from the analysis. We experimented with different decay times of one, three, and four years, and obtained very similar results. We re-introduced individuals if they formed a new tie in later years (e.g. Uzzi and Spiro, 2005).

We developed several definitions of a tie. Interactions in the mailing lists are structured in conversation threads (Orlikowski and Yates 1994). By conducting a content analysis of 200 randomly selected threads, we discovered that many responses to follow-ups in threads were directed to the individual that initiated the thread. To cope with this methodological challenge, we derived the networks in two different ways.

First, we created an affiliation, bipartite or two-mode network (Wasserman and Faust 1994) with two distinct sets of social entities: (1) the authors of the e-mails; and (2) the conversation threads in which these individuals interacted. We converted the two-mode data to one-mode data, representing interactions among individuals in one thread. When posting responses to threads, individuals are expected to have read earlier replies and are thus aware of others on a thread. The average size of a conversation thread in GNOME was only 3.05 messages (387,626 mails/127,046 conversation threads), so it is reasonable to assume that individuals interacting on the same threads are aware of each other, and that these therefore represent social ties. In other words, we assume that individuals interacting in the same thread have a tie and represent a fully connected clique (Newman, 2001). When converting to a one-mode data set, directionality is omitted and ties are treated as symmetric.

Second, we used the database of all interactions in the mailing lists to create directed networks. It could be the case that the 4th individual to interact on a conversation thread, may not be aware of all of the prior discussions in a conversation thread. To cope with this challenge, we constructed networks where we looked at who responded to whom to obtain directed ties. Using these networks, we derived in-degree centrality measuring the number of individuals that send an email directed to an individual i in year t. We also measured out-degree centrality that measures an individual i's interactions with other individuals in the community in year t.

The undirected network and the directed network measures were very correlated (>.87 for different degree centrality measures). In reading conversation threads, we discovered that many individuals were responding to the initial message of a thread but posted it to a follow-up. Because of this and the relative small size of 3.05 messages on average per thread, we used undirected ties. The alternative definitions of ties mentioned above yield the same substantive results.

(3) *Coordination Work*. Before the creation of the foundation, there were few formal mechanisms for coordinating project work. Two technical systems helped coordinate software development. The CVS system coordinated work by communicating code changes to everyone and the Bugzilla system coordinated work by showing the relevant tasks that were resolved and left unresolved. While these two systems helped project contributors manage interdependence within defined areas, there were few feedback mechanisms in place to foster the mutual adjustment needed to manage interdependence across subproject boundaries (March and Simon, 1958; Thompson, 1967). Thus, we measure coordination with boundary spanning behavior across sub-project communities and involvement in project-wide coordination discussions.

Number of sub-communities. We measured the *number of sub-communities*$_{it}$ or the extent to which an individual bridges sub-communities by the number of mailing lists an individual *i* has posted a message to in year *t*. With 137 active mailing lists in the community, not every individual has the time to join and attend to all mailing lists – anyone who tried would be overwhelmed. Each list is a sub-community focused on specialized interests. For example, one mailing list is devoted to discussions about a tool for profiling memory usage and finding memory leaks. Because posting the same message to multiple lists is discouraged, we are reasonably confident that our measure is reliable.

Number of coordination discussions initiated**.** By content coding the 137 mailing lists, we coded discussions devoted to coordination about project-wide issues. These included discussions focused on organizing project-wide affairs such as: code releases, conferences, the project website, or the foundation's governance. We developed a variable, the *number of coordination discussions initiated*$_{it}$ that measures the natural log of the number of new threads an individual *i* had initiated in year *t*. A very small percent of all communication within the community was devoted to coordination (never reaching 4% in any given year). Thus, obtaining significant results may constitute a very stringent test of coordination behavior.

Number of responses to coordination discussions**.** We developed another variable from these same data measuring the *Number of responses to coordination discussions*$_{it}$ for individual *i* in year *t*.

Control variables

Location**.** Individuals taking part in open source projects are often distributed across geographical boundaries. Yet, as shown by O'Mahony and Ferraro (2007), face-to-face interaction can be important to the governance of a community. We controlled for whether *location* influenced the possibility of gaining authority, by coding whether an individual was geographically located in the US. This was enabled by examining the time zone of sent mails (obviously controlling for discrepancies from the same individual) and the domain of the emails. We checked various secondary data on the GNOME web site to develop this measure and distinguish between individuals from South America versus those from North America.

Tenure**.** Since more experienced individuals were likely to have better knowledge about the latest developments; have better developed communication and collaboration networks and, by extension, be more knowledgeable about how to gain lateral authority (e.g. Pearce, 1993), we controlled for an individual's tenure. We therefore measured $tenure_{it}$ by the number of years of activity in the GNOME mailing lists for each individual *i* in year *t*.

Corporate affiliations**.** Some individuals in open source communities have affiliations with firms and participate in the community as part of their daily work (Dahlander and Wallin, 2006; O'Mahony and Bechky, 2008). Because these individuals are paid by firms to work on the project, they may be able to spend more time on the project which may affect their chances of progression. Using the information in emails, we coded whether an individual sent at least one mail from one of these firms in any given year. We identified two types of firm affiliation. First, *affiliation with an open source startup*$_{it}$ it i.e. a firm whose business model is based on GNOME (such as Helix Code, Red Hat, Eazel, Gnumatic and Henzai), and second, *affiliation with an incumbent in the software industry*$_{it}$ i.e. a large established firm participating in the GNOME community. These measures are coded for individual *i* in year *t* to control for any affects that might stem from an individual's corporate affiliation.

Estimation technique

We developed panel data with yearly observations for all individuals. We dealt with possible endogeneity concerns by using a one year lag for the independent variables in the specification of models to test whether network position at time *t*-1 influences the likelihood of becoming a new member or becoming elected at time *t*. Thus, the independent and control variables begin in 1999 (t-1), the year before the foundation was formed and members began to be appointed and board directors elected (2000).

Becoming a new member of the foundation is dichotomous with a very skewed distribution: there is a much larger share of zeros than 1s, as only a small proportion of community contributors become members. Logistic regressions assume that this distribution is symmetric, or that there are similar shares of zeros and 1s (Greene, 2000) which does not suit our data. To address this problem, we adopted a random-effects complementary log-log model, which fits maximum likelihood models with dichotomous dependent variables. Complementary log-log models are discrete-time survival models with a proportional hazard assumption (Wooldridge, 2002). For predicting elected board directors, the skewed distribution of zeros and 1s were less pronounced, and our tests showed we could use a random-effects logistic regression.

Hausman, Hall and Griliches (1984) advised testing to assess whether a random or fixed effects approach is most suitable for panel data. There are two reasons why a fixed effects model is not applicable. First, the complementary log-log model used in the analyses has no command for a conditional fixed-effects model, as there is no sufficient statistic that would allow the fixed effects to be conditioned out of the likelihood. Second, with the asymmetric distribution of zeros and 1s and the few individuals that actually do progress to membership or directorship, there is little variation in the dependent variable, which would result in many observations dropped from the analyses. In our regressions we include ρ values to show the panel-level, or the individuals' variance component. Following Greene's (2000) advice, when using random effects, we account for possible unobserved heterogeneity by including time invariant control variables.

RESULTS

Our descriptive data shows that the GNOME community was undergoing incredible growth at the time of our study. Table 1 shows that the number of individuals participating in GNOME almost doubled during the period of the study, reaching a peak in 2003 and levelling slightly in 2004. The total number of mails sent each year follows a similar trend, with a rapid increase from 38,428 in 1999 to a peak of 80,963 in 2003. With growth in the community's population, the number of active sub-communities increased dramatically – almost sixfold from 28 sub-community mailing lists in 1999 to 119 in 2004[2]. This increased specialization in community activities resulted in far less communication within each sub-community. The mean emails per mailing list decreased by almost half, from 1372.4 mails in 1999 to 638.9 mails in 2004. Thus, over the time of the study, the community became larger and increasingly specialized with far more sub-communities supporting far less activity. Perhaps it is not surprising that coordination discussions increased over time. However, coordination continued to remain a very small proportion of all communication within the community.

The number of commits, or contributions of source code to the project, increased during the first two years of the study with a peak of 381,602 commits in 2001. The number of contributions slightly declined thereafter, although the number of individuals committing code continued to rise. In 1999, 287 individuals committed code to the project, doubling to 577 individuals in 2004. Over the course of the study, 1137 individuals committed source code – far less than the number contributing to the mailing list (24,694). These data illustrate concentric circles of different types of community participation – with larger numbers contributing to project discussions, smaller numbers contributing code, and still smaller numbers yet achieving project membership or becoming elected to a board of director.

- Insert Table 1 about here -

[2] During the period under study, 137 mailing lists were active, but not all in the same year.

The antecedents of lateral authority

Our first seven hypotheses were directed toward predicting progression to positions of increasing lateral authority. Our data allow us to examine progression at two successive levels. Because each level of progression is a different dependent variable, we present the full set of results for the first seven hypotheses at each step of progression separately – moving from (1) becoming a member to (2) becoming elected to the board of directors. After presenting these results, we explore the consequences of gaining authority (H8) with differences-in-differences regressions at both levels of progression.

Progression Step 1: Becoming a member. Table 2 presents the descriptive statistics for the 24,694 individuals and the 74,249 individual years in the population.

- Insert Table 2 about here -

There are some high correlations between the independent variables. We therefore conducted several assessments to avoid problems of multicollinearity for both our complementary log-log and logistic regressions. As Menard (2002) notes, as the primary concern with multicollinearity is among the independent variables, the functional form of the dependent variable is less relevant. This suggests that one can adopt methods that are associated with OLS regressions, such as Variance Inflation Factors (VIFs). The highest VIF score was 3.5 and the average score 2.1, well below the generally accepted threshold level of 10 (Greene 2002). Second, we stepwise included the independent variables to assess how each one altered our results. The effects were stable across the different regressions, suggesting that multicollinearity was not a major issue (Menard 2002). Third, we randomly omitted 20% of the sample and created three sub-samples and reran the analyses. If multicollinearity was a major problem, coefficients should differ dramatically even with small changes in the sample (Greene, 2000). These sensitivity analyses provided strong support for our results. Given these three tests, we are confident that multicollinearity is not a major problem for our analyses.

Table 3 reports the results from the random-effects complementary log-log regressions, predicting *New Member* at time *t* as a function of the independent variables and control variables at time *t*-1. Model 1 is the baseline model including only the control variables. In Model 2 we include our variables measuring technical problem-solving. Model 3 adds the two variables measuring collaboration networks, and the final Model 4 includes our three coordination variables. We use the final Model 4 to interpret our results.

Technical problem solving is important in predicting membership. We find support for H1 as the positive and significant coefficient for *Number of Commits* indicates that the more code an individual has managed to get accepted by peer review, the more likely that individual is to gain lateral authority. This strong, significant effect indicates that making technical contributions that are accepted by the community is necessary for an individual to become a full fledged member of the community. Our measure of technical problem solving also considered an individual's ability to solve pre-existing problems as measured by the *Number of resolved bugs*. This effect is positive and significant, lending support for H2. Individuals that solve problems that are critical to the project as a whole are more likely to be granted membership in the community. We tested H3a and H3b by investigating an individual's contribution to technical discussions. The coefficient for the *Number of technical discussions initiated* is insignificant. We find no support for H3a, that the initiation of new technical discussions positively influences the likelihood of becoming a member. However, lending support to H3b, the coefficient for *Number of responses to technical discussions* is positive and significant. Responding to discussions about technical issues is more likely to be rewarded with membership in the community than initiating technical discussions.

Collaboration networks are important to predicting membership – but these collaborations must be meaningful technical collaborations as opposed to merely being central in the project's communication networks. The coefficient for the *Number of technical collaborations* is positive and significant, so we do find support for H4 – the number of technical collaborations an individual engages in leads to a higher likelihood of gaining membership on the project. However, the coefficient for *Structurally central in communication network* is insignificant. The likelihood of gaining lateral authority was unaffected by an individual's sheer number of ties with other individuals in the community.

All three of our coordination hypotheses are supported. The coefficient for *Sub-community participation* is positive and significant, lending support for H6. Community contributors that engage in dialogue across more sub-communities will be more likely to gain membership. The positive and significant coefficients for the *Number of coordination discussions initiated* indicates that creating new dialogue about coordination issues positively impacts the likelihood of gaining membership – supporting H7a. The coefficient is positive and significant at the 1% level, despite the fact that coordination discussions are such a small percentage of all project communication. Likewise, the *Number of responses to coordination discussions* has a positive and significant coefficient, also at the 1% level, lending support to H7b. The more an individual responds to coordination discussions initiated by others, the more likely an individual is to become a member of the community.

- Insert Table 3 about here -

Step 2: Becoming elected to the board of directors. Table 4 presents the descriptive statistics for the 616 individuals and the 1,619 individual years that are members of the foundation at year *t*. Since only current members can be elected to the project's board of directors, the appropriate population of study is foundation members as opposed to all individuals who have contributed to the project.[3]

- Insert Table 4 about here -

Table 5 presents random-effects logistic regression estimates of individuals' likelihood of becoming elected to the board of directors at time *t*, as a function of the independent variables and control variables at time *t*-1. Model 5 is the baseline and only includes the control variables. In Models 6-8, we stepwise include the independent variables.

- Insert Table 5 about here -

Unlike our regression results predicting membership, the results reported in Table 5 provide no support for technical problem solving as an important predictor of progression to board director. We find no support for H1 proposing that the *Number of commits* influences the likelihood of becoming an elected board director. The coefficient for the *Number of resolved bugs* is also insignificant, which leads us to reject H2 at the second stage of progression. We also find that the coefficients for *Number of technical discussions initiated* and *Number of responses to technical discussions* are insignificant at this stage, failing to support H3a and

[3] It could be the case that the same variables that affects an individual's likelihood of becoming a member drives becoming a board director. In alternative estimations we used a Heckman (1979) correction procedure that yields the same results. We chose not to report these for space considerations.

H3b. While technical problem solving is important for predicting progression to membership, we find no support for H1-3 predicting the importance of technical contribution for progression to the board of directors.

With respect to collaboration networks, we find different results for our two measures. The coefficient for *Structurally central in communication network* is positive and significant at the 1% level, lending support to H4: individuals that are structurally central in the community are more likely to become elected to the board. However, we do not find support for H5 as the coefficient for the *Number of technical collaborations* in code is insignificant. Individuals with more technical collaborations are not more likely to become elected board directors. In direct contrast with our results predicting membership, what matters for the board of directors is not technical collaborations on code but centrality in a project's communication networks.

Table 5 provide strong support for all of our coordination hypotheses at the board director level. The *Sub-community participation* coefficient is positive and significant at the 1% level, supporting H6's prediction of a positive effect of engaging in several sub-communities on the likelihood of becoming a board director. Individuals that span more sub-community boundaries are more likely to become elected to the board. We also find support for H7a and H7b that proposed a positive relationship between initiating and responding to new coordination discussions and the likelihood of gaining lateral authority. Both the coefficients for the *Number of coordination discussions initiated* and the *Number of responses to coordination discussions* are significant at the 5% level.

While all of our coordination hypotheses are supported at both levels of progression, taken together, our results suggest that different mechanisms influence progression at different levels. To become a member, it is important to make technical contributions that are accepted by peers in the community; to resolve technical problems; and respond to technical discussions within the community. To become elected to the board of directors, however, the results are very different. Spanning sub-community boundaries, establishing and maintaining social ties and initiating discussions and responding to coordination discussions are critical talents needed to garner the community's confidence in a board director.

The consequences of lateral authority

We now turn our analyses to assess the potential consequences of lateral authority. Hypothesis 8 argued that after individuals gained lateral authority they would increase their coordination work on the project. To disentangle this question, we examined whether our independent variables changed before and after individuals moved to each step of authority. More specifically, we examine the individuals that become a member or director by comparing their behavior the year before this event with their behavior in the immediate year following. In order to meaningfully compare individual behavior before and after becoming a new member, we focus on those observations where we have data on the subsequent year (some individuals drop out of the analysis as they are not observed after 2004). This analysis is based on 547 individuals that we compare for two years. However, individuals that gained authority could have changed their behaviour due to other factors at the time of the treatment. Thus we use a differences-in-differences approach to compare a treatment group (that gained authority) with a control group (that did not). We matched individuals that became members with other individuals in the community that did not, but still shared similar characteristics. The total sample therefore equals 547*2 *2=2188 observations. See Appendix A for a complete description of the matching and the analytical technique. With an identified control group, we can isolate the treatment effect associated with gaining authority (Bertrand et al., 2004). Table 6 shows the results from the differences-in-differences regressions.

- Insert Table 6 about here -

To conserve space, we report whether the differences-in-differences coefficient was positive/negative at the 5% level in Table 6 for both members and board directors. A positive sign indicates that individuals that become members or board directors significantly increased their activity within a year of assuming the position. We find that there is an immediate and significant change in behavior within a year of an individual gaining lateral authority relative to the control group. The first column of Table 6 reports the results from differences-in-differences regressions predicting the effects of membership. Indeed, when individuals gain membership they become *increasingly* active across all three categories (technical problem solving, collaboration networks, and coordination work). Gaining more authority over the project - even at the level of membership results in intensification of effort to the project.

In the second column of Table 6, we report the results from differences-in-differences regressions predicting individual's change in behavior after becoming elected to the board of directors. This analysis is based on 28 individuals before and after being elected as well as the control group of the same size (28*2*2). This sample is therefore smaller than the one for members. Once individuals join the board of directors, they significantly reduce their technical efforts, contributing less source code and resolving fewer bugs. After progressing to the board of directors, individuals dramatically increase their coordination work – increasing the number of coordination discussions they initiate and respond to as well as increasing their boundary spanning behavior across sub-communities.

Together, these results suggest that gaining lateral authority has a dramatic effect on individual contributions to the project – encouraging much more activity than those same individuals exhibited only one year prior to progression. Although only members increase their technical and collaboration efforts, the effects for coordination behavior are strong across both levels of progression. Because our preliminary field data suggested that our informants were blasé about the authority positions they acquired; quick to downplay the importance of authority; and demonstrated some reluctance to assume a position of authority in the first place, we were surprised by the magnitude of the changes in behavior that we found. We suspect that lateral authority gives individuals greater license to help coordinate the work of their peers and integrate the increasingly specialized sub-communities that formed in the project over time. For those who argue that open source community projects are "self managed", it may seem contradictory to discover that clear authority roles enable coordination work, yet this discovery does much to further our understanding of how community managed projects function.

DISCUSSION

Organization scholars have noted that when work is more likely to require technical skill or the manipulation of information (Barley and Orr, 1997), horizontal forms of organizing would become increasingly likely (Barley and Kunda, 2006). High trust, relational ties, and adaptable structures – all features of community and project forms, are theorized to foster the generation and exchange of knowledge and know-how (Powell, 1990; Adler and Borys, 1996; Adler, 2001; Adler and Heckscher, 2006). However, few have examined the unique problems of authority and coordination that can occur within projects that are not situated within a single organization – where members may not have direct authority over those upon whom they depend upon to create new knowledge, or in this case – develop software. According to Simon, a shared authority system "distinguishes the behavior of individuals as participants of organizations from their behavior outside such organizations." (1976: 124). Simon acknowledged that "authority…can operate "upward" and sidewise" as well as "downward" in the organization" (1976: 12), but he did not elaborate how lateral forms of authority might differ from vertical lines of authority.

The literature on scientific and technical careers did make a case for the need for horizontal modes of progression or careers of achievement (Zabusky and Barley, 1996) but primarily focused on the conflict that can occur when highly technical or specialized experts progress along a 'separate but equal' hierarchy. Our approach was to examine how progression occurred within a peer managed system where no such conflict existed. The open source community project we studied is an example of an organizational form where "there are gains to be had by pooling resources" (Powell, 1990: 303) and individuals contribute to obtain both individual and mutual benefits (von Krogh and von Hippel, 2006). Yet, as our research shows, opportunities to 'progress to the center' and gain lateral authority over collective work still exist.

Scholars have been slow to appreciate how community forms coordinate their efforts without formal authority structures (Adler, 2001; Adler and Heckscher, 2006). Our informants revealed that this could be a frustrating experience. The GNOME community's approach to this problem was to create two progressive stages of lateral authority for aspiring individuals. However, which individuals would gain access to this system was not clear – even to the designers of the system. The theoretical question we explored is - what predicts individual progression in this context? And, of equal importance, how does progression to new authority roles affect individuals' subsequent behavior?

Drawing upon the literature on knowledge work in communities, we identified three types of behaviors that could predict progression in community forms: (1) technical problem solving; (2) collaboration networks; and (3) coordination work. We then tested these different explanations on a longitudinal data set that included two successive levels of progression to explicate the antecedents of acquiring lateral authority. The ability to evaluate how cohorts move through successive stages of progression over time is rare in network research (Burt, 2000) and allows us to parse how mechanisms differ by stage.

We found that engaging in solving technical problems and technical collaborations with others explained why some individuals became members, but was less important in predicting progression to the board of directors. Technical problem solving can be seen as a "ticket of admission" to gaining lateral authority, but the effect disappears at the next stage. These membership results will not surprise scholars who have argued that problem solving and collaboration behaviors are some of the long-standing benefits associated with participation in a technical community (Brown and Duguid 1991; 2001). However, our two stage model provides a better understanding of how mechanisms change as individuals progress toward the center – and at later stages coordination becomes more important. Most interestingly, all of our hypotheses with regard to coordination were important to predicting progression at both member and board of director levels. Despite the fact that coordination emails accounted for a small percentage of all communications in the community (no more than 3.5 percent in any given year), involvement in such discussions had positive and significant effects.

Finally, we were curious to know how the experience of gaining authority affected subsequent individual behavior. Using a differences-in-differences approach, we found strong evidence that even within a short period, individuals significantly increased their contributions to the project after gaining a position of authority - despite our informants' claims to the contrary. For new members, these effects were across the board. For new members of the board, technical problem solving actually declined while engagement in coordination work increased significantly. Membership seems to deepen all individual commitment and effort to the project, while becoming a board director triggers only more coordination work. Our findings offer a striking picture of individuals stepping in to take on coordination work – only after they have been given the license to do so. In a community where "no one is in charge", individuals may be reluctant to engage in essential coordination work for fear of violating cultural norms that value individual autonomy (e.g. O'Mahony and Chen, 2008; O'Mahony, 2007). When interpreting these results in light of the community's growing population and increasingly fragmented sub-community structure over our period of study, our findings suggest that the effects of lateral authority offer an important integrating mechanism that may have been previously missing.

Our investigation of the antecedents and consequences of lateral authority in a mature open source community enable us to make three distinct contributions to organizational theory: (1) a more specific conception of lateral authority: how it works in projects organized in communities and the consequences for progression; (2) an understanding of the different mechanisms that affect successive levels of progression; and (3) theoretical insights as to how lateral authority affects the coordination of knowledge work.

Specifying Lateral Authority and the Consequences for Progression. Despite the fact that lateral authority has been proposed as a critical, distinguishing feature of several types of organizing forms that fill the growing void between market and hierarchy – namely, community (Adler, 2001; van Maanen and Barley, 1994), project (Bechky, 2006; Jones & Lichenstein, 2008), and network forms (Powell, 1990), little attention has been devoted to how lateral authority differs from our prior conceptions of vertical authority and how this might affect traditional notions of progression. Since Etzioni (1959), it has been clear that it is difficult for any organization to realize its goals without a clear authority structure. As Harrison (1960:233) put it: "No social system can operate on a continuous basis without support from some mode of authority, no matter how informal and inadequately defined it may be". Yet, few empirical studies have explored lateral authority in any depth. Coleman theorized that one reason for the demise of interest in this line of work was scholars' static interpretations of authority (1980: 152). Our research presents a more dynamic conception of lateral authority by showing the conditions and limitations under which it was granted; the behaviors that predict which individuals are more likely to acquire it and the consequences for individuals and projects.

Scholars of technical and scientific workers were early to understand the future of knowledge work when they recognized that in order to retain, attract and develop technical talent, that industry would need to move toward a system of 'colleague authority'" Marcson (1960: 72). However, by focusing on the dyadic "strain and conflict" that unfolded *between* managers and industrial researchers as opposed to how researchers *themselves* might coordinate their work, this literature failed to consider how lateral authority might affect progression patterns.

Studies of technical experts have long argued that talented individuals dedicated to their discipline or craft often do not desire 'vertical' careers but instead careers of achievement (Zabusy, 1997; Zabusky and Barley, 1996; Arthur, 1994; Arthur and Rousseau, 1996) but this literature has not previously been well connected to our understanding of community and project forms. Our research shows how a peer based community of individuals created a system to transfer authority over their collective project to specific individuals – but without establishing a traditional vertical hierarchy. By crafting a lateral basis of authority, a new conception of progression emerges whereby dedicated individuals can progress toward the center – gaining authority over the project without authority over individuals. This in turn has consequences for the coordination of project work as well as for individual careers. The phenomena of gaining lateral authority and progressing to the center is not likely to be relevant just to open source projects, but apply to many contexts where experts and specialists of all kinds collaborate to achieve collective outcomes (Hargadon and Becky, 2006).

Mechanisms Affecting Progression. By combining several different sources of data to study cohorts progressing over a six-year period, we were able to parse how the mechanisms affecting progression changed at two distinct levels. We show clearly that technical contributions and collaborations matter only to acquiring membership and that beyond this point, coordination work and centrality in the project's communication network are critical to acquiring elected board positions. The fact that individuals in this context do less hands on technical work and more coordination work as they move toward the center is not entirely inconsistent with earlier research on the essence of managerial work (Mintzberg, 1973).

Yet, the behavioural predictors of progression on a community managed project that emerge as most consistently important put coordination work (by any measure) front and center and this is despite the fact that coordination work was a small percent of all community activity. Mintzberg gave short shrift to coordination work and our data suggest that the challenge of balancing differentiation and integration first identified by Lawrence and Lorsch (1967) remains salient on a community managed project. The problem is there are few mechanisms available to help balance this tension. Thus boundary spanning work and engagement in coordination discussions becomes more important. By exposing the importance of coordination work as an antecedent to authority, we contribute an understanding of how individuals contributing to complex knowledge producing enterprises actually “self-organize” their collective work. As opposed to managing individuals, what becomes essential is the coordination and integration of different subsystems of the project.

In terms of mechanisms, what may surprise network theorists (Podolny and Baron, 1997), is that an individual’s structural centrality was not predictive of the first stage of progression - membership. From a social network perspective, centrality in the project’s communication network should be important to achieving both levels of progression as both are determined by peers. However, we have two different populations in the pipeline at each stage – all contributors predicting the member stage and the smaller subset of members predicting the board of director stage. What may explain these results is that a member’s centrality in the project is more salient and visible than the centrality of someone who has only contributed to the project. Thus, membership may affect the way an individual's centrality in the network is perceived (Krackhardt, 1990).

Coordinating Knowledge Work. Scholars examining coordination practices often look at the mechanisms that allow individuals in either distributed or collocated locations to coordinate individual work efforts into an integrated whole (Okhuysen and Bechky, 2008). Organizing structures (March and Simon, 1958; Ouchi, 1980), boundary objects (Bechky, 2003a,b; Carlile, 2002), schedules, plans, meetings (Gittell, 2002), and routines (Feldman, 2000; Nelson and Winter, 1982) all help individuals manage the dual processes that constitute coordination: disaggregating work into smaller parts that can be accomplished by individuals and reintegrating those contributions into a collective whole (Okyhusen and Bechky, 2008; Lawrence and Lorsch, 1967).

In a community where contributors are distributed around the world, many traditional forms of coordination devices such as face-to-face meetings and physical boundary objects are not available. Furthermore, the culture of open source projects explicitly rejected routines, schedules, and anything reeking of 'too much structure' as an anathema to their ethos of free, unfettered software development (O'Mahony and Bechky, 2008; Chen and O'Mahony, 2008). When the community began to scale, organizers developed more sub-communities to focus individual attention on specific parts of the project (e.g. such as sound, video, mobile, security) but specialization led to reduced communication within each sub-community. Few mechanisms ensured that sub-communities would continue to communicate with each other – an essential condition for coordinating technological work (Sosa, Eppinger, Rowles, 2007). Thus, differentiation occurred without any corresponding integration mechanisms – a danger that Lawrence and Lorsch (1967) warned against.

Perhaps it is not a surprise to see individuals accelerating their coordination work after acquiring a position of lateral authority on the project. However, given our understanding of the community's culture and norms, we did not expect to see authority have such a dramatic effect. Several mechanisms may be responsible for these somewhat surprising results. First, an access to information argument: after moving into their new role, members and board directors may possess more information that may prompt them to engage with a broader number of sub-communities in order to diffuse that information and help coordinate the community's efforts. Second, a role expectation argument: once in the role of member or board director, individuals may feel that there is an expectation for them to assume a more visible role in the community and thus engage in more boundary spanning and coordination work – with this effect being greater for board of directors. Finally, a legitimacy argument: when individuals gain some degree of authority, the legitimacy that comes with this role may provide them with the license to engage in more coordination behavior. In a community where norms against hierarchy are strong and coordination mechanisms weak, it may be difficult for individuals to take on additional coordination tasks even if there is a need for them. All three explanations address one underlying mechanism: role clarity. Indeed, our research suggests somewhat paradoxically that the specification and granting of authority roles may foster coordination work - particularly in high growth settings where these parameters are previously unspecified.

While boundary spanning has long been viewed as important to individual and team success in organizations and to supporting innovation outcomes (Tushman, 1977; Ancona and Caldwell, 1992; Obstfeld, 2005; Fleming and Waguespack, 2007; Fleming, Mingo, Chen, 2007), the link between boundary spanning activities and coordination work has been only loosely specified. Our research suggests that this link may be more critical than previously acknowledged. We encourage scholars to examine this relationship more specifically as it is likely to be an important predictor of the sustenance or success of collective work – especially when that work is distributed and spans organizations.

REFERENCES

Adler, P. S.
2001 "Market, hierarchy, and trust: the knowledge economy and the future of capitalism." Organization Science, 12(2): 214-234.

Adler, P.S. and C. Heckscher.
2006 "Towards collaborative community" in C. Heckscher and P. S. Adler (Eds.), The Firm as a Collaborative Community, New York: Oxford University Press.

Adler, P.S. and B. Borys.
1996 "Two types of bureaucracy: Coercive versus enabling." Administrative Science Quarterly, 41(1): 61-89.

Allen, T.
1977 Managing the Flow of Technology: Technology Transfer and the Dissemination of Technological Information Within the R and D Organisation. Cambridge, Mass.: The MIT Press.

Allen, T. and R. Katz.
1986. "The dual ladder: Motivational solution or managerial delusion?" R&D Management, 16(2): 185-197.

Althauser, R. P.
1989 "Internal labor markets." Annual Review of Sociology, 15: 143-161.

Ancona, D. G. and D. F. Caldwell.
1992 "Bridging the boundary External activity and performance in organizational teams." Administrative Science Quarterly, 37(4): 634-665.

Arthur, M. B. and D. M. Rousseau.
1996 The Boundaryless Career: A New Employment Principle for a New Organizational Era. Oxford: Oxford University Press.

Arthur, M. B.
1994 "The boundaryless career: A new perspective for organizational inquiry." Journal of Organizational Behavior, 5: 295-306.

Bailyn, L.
1991 "The hybrid career: An exploratory study of career routes in R&D." Journal of Engineering and Technology Management, 8: 1-14.

Baldwin, C. Y. and K. B. Clark.
2006 "The architecture of participation: Does code architecture mitigate free riding in the open source development model?" Management Science, 52: 1116-1127.

Barley, S., R.
1995 "The New World of Work. Washington: British North American Committee.

Barley, S.R. and J. E. Orr.

1997 Between Craft and Science: Technical Work in U.S. Settings. Ithaca, NY: ILR Press.

Barley, S. R. and G. Kunda
2001 "Bringing work back in." Organization Science, 12: 76-95.
2004 Gurus, Hired Guns, and Warm Bodies: Itinerant Experts in a Knowledge Economy. Princeton, NJ: Princeton University Press.

Bechky, B. A.
2003a "Object lessons: Workplace artifacts as representations of occupational jurisdiction." American Journal of Sociology, 109: 720-52.
2003b "Sharing meaning across occupational communities: The transformation of knowledge on a production floor." Organization Science, 14: 312-330.
2006 "Gaffers, gofers, and grips: Role-based coordination in temporary organizations." Organization Science, 17: 3-21.

Bertrand, M. E. Duflo, and S. Mullainathan.
2004 "How much should we trust differences-in-differences estimates?" Quarterly Journal of Economics, February: 249-275.

Bourdieu, P.
1990 The Logic of Practice. Stanford, CA: Stanford University Press.

Brass, D. J.
1984 "Being in the right place: A structural analysis of individual influence in an organization." Administrative Science Quarterly, 29: 518-539.

Brown, J. S. and P. Duguid.
1991 "Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation." Organization Science, 2: 40-57.
2001 "Knowledge and organization: A social-practice perspective." Organization Science, 12: 198-213.

Burt, R. S.
1992 Structural Holes: The Social Structure of Competition. Cambridge, MA.
2000 "Decay functions." Social Networks, 22: 1.

Butler, B.
2001 "Membership size, communication activity and sustainability: A resource-based model of online social structure." Information Systems Research, 12: 346-362.

Carlile, P. R.
2002 "A pragmatic view of knowledge and boundaries: Boundary objects in new product development." Organization Science, 13: 442-455.

Cassell, J., D. Huffaker and D. Tversky.
2005 How to win a world election: Emergent leadership in an international online community. Paper presented at the Proceedings of the Second Communities and Technologies Conference, Milan.

Coleman, J. S.

1980 "Authority systems." Public Opinion Quarterly, 44: 143–163.

Crowston, K.
1997 "A coordination theory approach to organizational process design." Organization Science, 8(2), 157-175.

Dahlander, L. and M. W. Wallin
2006 "A man on the inside: Unlocking communities as complementary assets." Research Policy, 35: 1243-1259.

Daft, R. L. and A. Y. Lewin.
1993 "Where are the theories for the new organizational forms - an editorial essay." Organization Science, 4: R1-R6.

Doeringer, P. and M. Piore.
1971 Internal Labor Markets and Manpower Analysis. Lexington, Mass: Heath.

Etzioni, A.
1959 "Authority structure and organizational effectiveness." Administrative Science Quarterly, 4: 43-67.

Feldman, M. S.
2000. "Organizational routines as a source of continuous change." Organization Science, 11: 611-629.

Fleming, L. and D. M. Waguespack.
2007 "Brokerage, boundary spanning, and leadership in open innovation communities." Organization Science, 18(2): 165-180.

Fleming, L, S. Mingo, and D. Chen.
2007 "Collaborative brokerage, generative creativity, and creative success." Administrative Science Quarterly, 52(3): 443-475.

Freeman, L. C.
1979 "Centrality in social networks." Social Networks, 1: 215-240.

Gittell, J. H.
2000 "Paradox of coordination and control". California Management Review, 42(3): 177-183.

Gittell, J. H.
2002 "Coordinating mechanisms in care provider groups: Relational coordination as a mediator and input uncertainty as a moderator of performance effects. Management Science, 48: 1408-1426.

Goldner, F. H. and R. R. Ritti.
1967 "Professionalization as career immobility." American Journal of Sociology, 72(5): 489-502.

Gould, R. V.

2002 "The origins of status hierarchies: A formal theory and empirical test." American Journal of Sociology, 107: 1143-1178.

Gouldner, A. W.
1968 "Cosmopolitans and locals: Toward an analysis of latent social roles." Administrative Science Quarterly, 2: 281-306.

Greene, W. H.
2000 Econometric Analysis, 4th edition. Upper Saddle River, New Jersey: Prentice-Hall.

Hansen, M.
1999 "The search-transfer process problem: the role of weak ties in sharing knowledge across organization subunits." Administrative Science Quarterly, 44: 80-110.

Harrison, P. M.
1960 "Weber's categories of authority and voluntary associations." American Sociological Review, 25: 232-237.

Hargadon, A. and R. Sutton.
1997 "Technology brokering and innovation in a product development firm." Administrative Science Quarterly, 42: 716-749.

Hargadon, A. B. and B. A. Bechky.
2006. "When collections of creatives become creative collectives: A field study of problem solving at work." Organization Science, 17: 484-500.

Hausman, J. , B. Hall, and Z. Griliches.
1984 "Econometric models for count data with an application to the patent - R&D relationship." Ecometrica, 52: 909-938.

Heckman, J. J.
1979 "Sample selection bias as a specification error." Econometrica 47(1) 153-161.

Ibarra, H.
2004 Working identity: Unconventional strategies for reinventing your career. Boston: Harvard Business School Press.

Jones, C., W. S. Hesterly, and S. P. Borgatti
1997 "A general theory of network governance: Exchange conditions and social mechanisms." Academy of Management Review, 22: 911-945.

Jones, C. and B. B. Lichtenstein.
2008 "Temporary Interorganizational Projects: How Temporal and Social Embeddedness Enhance Coordination and Manage Uncertainty". In Cropper, S., Ebers, M., Ring, P.S. & Huxham, C. (Eds.), Handbook of Interorganizational Relations: 231-255. London: Oxford University Press.

Katz, R. and M. L. Tushman.

1983. “A longitudinal study of the effects of boundary spanning supervision on turnover and promotion in research and development.” Academy of Management Journal, 26:437-456.

Katz, R., M. Tushman, and T. Allen.
1995 “The influence of supervisory promotion and network location on subordinate careers in a dual ladder R&D setting.” Management Science, 41(5): 848-863.

Kogut, B. and A. Metiu.
2001 “Open-source software development and distributed innovation.” Oxford Review of Economic Policy, 14: 248-264.

Kornhauser, W.
1962 Scientists in Industry: Conflict and Accommodation, Berkeley, CA: University of California Press.

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

Kuk, G.
2006 “Strategic interaction and knowledge sharing in the KDE developer mailing list.” Management Science, 52: 1031-1042.

Lakhani, K. R. and E. von Hippel.
2003 “How open source software works: "free" user-to-user assistance.” Research Policy, 32(6): 923-943.

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

Lam, A.
2005. “Work roles and careers of R&D scientists in network organizations.” Industrial Relations, 44(2): 242-275.

Lave, J. C. and E. C. Wenger.
1991 Situated Learning: Legitimate Peripheral Participation. New York: Cambridge University Press.

Lawrence, P. and J. Lorsch.
1967. “ Differentiation and integration in complex organizations.” Administrative Science Quarterly, 12 :1-30.

Lee, G. K. and R. E. Cole.
2003 “From a firm-based to a community-based model of knowledge creation: The case of the Linux kernel development.” Organization Science, 14: 633-649.

Long-Lingo, E. and S. O’Mahony.

2008 "Nexus work: Managing ambiguity on creative projects." Working Paper presented at: " Creativity, Entrepreneurship, and Organizations of the Future," a Harvard Business School Centennial Colloquium

MacCormack, A., J. Rusnak, and C. Y. Baldwin.
2006 "Exploring the structure of complex software designs: An empirical study of open source and proprietary code." Management Science 52(7): 1015-1031.

Malone, T. W. and K. Crowston.
1994 "The interdisciplinary theory of coordination." ACM Computing Surveys, 26(1):87-119.

Marcson, S.
1960 The Scientist in American Industry. Princeton: Princeton University.

March, J. G. and H. A. Simon.
1958 Organizations. NY: Blackwell.

Menard, S. W.
2002 Applied Logistic Regression Analysis. Thousand Oaks, CA: Sage University Paper Series.

McDaniel, S., G. Olson, and J. Magee.
1996 "Identifying and analyzing multiple threads in computer-mediated and face to face conversations". In Proceedings of the Conference on Computer Supported Cooperative Work: 39-47. Cambridge, MA: ACM Press.

Miles, R. E. and C. C. Snow.
1986 "Organizations: New concepts for new forms." California Management Review, XXVIII (3): 62-73.

Mintzberg, H.
1973 The Nature of Managerial Work. New York: Harper & Row.

Misiolek, N., and R. Heckman.
2005 Patterns of emergent leadership in virtual teams. Paper presented at the Proceedings of the 38th Hawaii International Conference on System Sciences, Hawaii.

Newman, M. E. J.
2001 "The structure of scientific collaboration networks." Proceeding of the National Academy of Sciences, 98: 404-409.

Obstfeld, D.
2005 "Social networks, the tertius iungens orientation, and involvement in innovation." Administrative Science Quarterly, 50: 100-130.

Okhuysen, G. and B. A. Bechky.
2008 Coordination in organizations: A relational perspective. Forthcoming Annals of Academy of Management.

O’Mahony, S.
2007 “What does it mean to be community managed?” Journal of Management and Governance, 11(2): 139-150.

O'Mahony, S. and B. A. Bechky.
2006 “Stretchwork: Managing the career progression paradox in external labor markets.” Academy of Management Journal, 49(5): 918-941.

O’Mahony, S. and F. Ferraro.
2007 “Governance in collective production communities.” Academy of Management Journal, 50(5): 1079-1106.

Ouchi, W. G.
1980 “Markets, bureaucracies and clans.” Administrative Science Quarterly, 23: 293-317.

Orlikowski, W. J. and J. Yates
1994 “Genre repertoire: The structuring of communicative practices in organizations.” Administrative Science Quarterly, 39: 541-574.

Osterman, P.
1984 Internal Labor Markets. Cambridge: MIT Press.

Pearce, J. L.
1993 Volunteers: The organizational behavior of unpaid workers. London: Routledge.

Pinchot, E. and G. Pinchot.
1993 The End of Bureaucracy and the Rise of the Intelligent Organization. Berrett-Koehler Publishers

Podolny, J. M. and J. N. Baron.
1997 “Resources and relationships: Social networks and mobility in the workplace.” American Sociological Review, 62: 673-693.

Powell, W.W.
1990 Neither market nor hierarchy: Network forms of organization. In B.M. Staw & L.L. Cummings (Eds.), Research in Organizational Behavior, 12, 295-336

Powell, W. W., K. W. Koput, and L. Smith-Doerr.
1996 “Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology.” Administrative Science Quarterly, 41: 116-145.

Powell, W. W., D. R. White, K. W. Koput and J. Owen-Smith.
2005 “Network dynamics and field evolution: The growth of inter-organizational collaboration in the life sciences”. American Journal of Sociology, 110(4): 1132-1205.

Powell, W. W. and K. Snellman.
2004 “The knowledge economy.” Annual Review of Sociology, 30: 199-220.

Raymond, E. S.

1999 The Cathedral and the Bazaar, Second edition. Sebastapol, CA: O'Reilly.

Romanelli, E.
1991 "The evolution of new organizational forms." Annual Review of Sociology, 17: 79-103.

Rosenkopf, L., A. Metiu, and V. George.
2001 "From the bottom up? Technical committee activity and alliance formation." Administrative Science Quarterly, 46: 748-772.

Ritti, R.
1968 Work Goals of Scientists and Engineers. Berkeley: Institute of Industrial Relations.
1971 The Engineer and the Industrial Corporation. New York, NY: Columbia University Press.

Rothman, R. A. and R. Perrucci.
1970 "Organizational careers and professional expertise." Administrative Science Quarterly, 15(3): 282-293.

Simon, H.
1976 Administrative Behavior, New York: Free Press.

Soda, G., A. Usai, and A. Zaheer
2004 "Network memory: The influence of past and current networks on performance." Academy of Management Journal, 47: 893-906.

Sosa, M.E., Eppinger, S. D. and Rowles, C. M.
2007 "Are your engineers talking to one another when they should?" Harvard Business Review, November 1.

Sproull, L. and S. Kiesler
1991 Connections: New Ways of Working in the Networked Organization. Cambridge, MA: MIT Press.

Thompson, J. D.
1967 Organizations in Action. New York: McGraw-Hill.

Turner, J. H.
1986 The Structure of Sociological Theory. Chicago, IL: University of Chicago Press.

Tushman, M. L.
1977 "Special boundary roles in the innovation process." Administrative Science Quarterly, 22: 587-605.

Uzzi, B. and J. Spiro.
2005 "Collaboration and creativity: The small world problem." American Journal of Sociology, 111: 447-504.

Van Maanen, J. and S. R. Barley

1984 "Occupational communities: Culture and control in organizations". Research in Organizational Behavior, 6: 287-365.

von Hippel, E.
2005 Democratizing Innovation, Cambridge, MA: MIT Press.

von Hippel, E. and G. von Krogh.
2003 "Open source software and the 'private-collective' innovation model: Issues for organization science." Organization Science, 14: 209-223.

von Krogh, G., S. Spaeth, and K. Lakhani.
2003 "Community, joining, and specialization in open source software innovation: A case study." Research Policy, 32(7): 1217-1241.

von Krogh, G. and E. von Hippel.
2006 "The promise of research on open source software." Management Science, 52: 975-983.

Wasserman, S. and K. Faust.
1994 Social Networks: Methods and Applications. Cambridge: Cambridge University Press.

Weick, K.
1993 "The collapse of sensemaking in organizations: The Mann Gulch disaster." Administrative Science Quarterly, 38: 628-652.

Weick, K. and K. Roberts.
1993 "Collective mind in organizations: Heedful interrelating on flight decks." Administrative Science Quarterly, 38(3): 357-381.

Wenger, E.
1998 Communities of Practice: Learning, Meaning, and Identity. Cambridge, UK: Cambridge University Press.
2000 "Communities of practice and social learning systems." Organization, 7 (2): 225-246.

Whalley, P. and S. R. Barley
1997 "Technical work in the division of labor: Stalking the wily anomaly." In S. R. Barley, and J. E. Orr (Eds.), Beyond craft and science: 23-52. Ithaca: Cornell University Press.

Wooldridge, J. M.
2002 Econometric Analysis of Cross Section and Panel Data. Cambridge: MIT Press.

Yates, J., W. Orlikowski, and S. L. Woerner.
2003 "Virtual organizing: Using threads to coordinate distributed work.", IEEE Proceedings of the Hawai'i International Conference on System Sciences, January 6-9, 21003.

Yoo, Y. and M. Alavi.

2004 “Emergent leadership in virtual teams: What do emergent leaders do?” Information and Organization, 14: 27-28.

Zabusky, S.
1997 “Computers, clients, and expertise: Negotiating technical identities in a nontechnical world.” in Barley, S.R. and J. E. Orr (Eds.) Between Craft and Science: Technical Work in U.S. Settings: 129-153. Ithaca, NY: ILR Press.

Zabusky, S. E. and Barley, S. R.
1996 “Redefining success: Ethnographic observations on the careers of technicians." in Osterman, P. (Ed.) Broken Ladders: White Collar Careers: 185-214. Oxford: Oxford University Press.