SOCIAL FORCES AND CONSTRAINT IN THE ATTAINMENT OF COMMUNITY STATUS*

Daniel Stewart
Washington State University

WORD COUNT:
13,998

* Direct all correspondence to Daniel Stewart, Washington State University, Ad Annex 203, PO Box 1495, Spokane, WA 99210-1495 (danstewart@wsu.edu). I am grateful to Bill Barnett for his advice and consistent encouragement on this project. Jim Baron, Joel Podolny, and Mikolaj Piskorski also contributed tremendously helpful insights that helped refine the analysis, as did seminar participants at Stanford University.

SOCIAL FORCES AND CONSTRAINT IN THE ATTAINMENT OF COMMUNITY STATUS

ABSTRACT:

Although there has been a fair amount of conjecture regarding the circumstances that lead to the generation and stability of status orders, most of the previous work in this area has typically studied the effects of stability by assuming that a status order has evolved endogenously, but without specifying empirically the underlying causes that create rigidity within the system. This paper investigates the manner in which social forces exert pressure that initially shapes and defines an actor's status within a community, but eventually constrain that actor's movement within a status order. The results of empirical analyses at the dyad level show that, in the process of status attainment, community members will tend to evaluate a focal actor's reputation according to publicly available social cues. Ironically, these same social cues eventually work to produce stability and constraint in an actor's status position by reducing heterogeneity in community status beliefs. Thus, as the cumulative number and average status of others making references about a focal actor's status increases, the likelihood of that actor being recognized as holding a status that differs from these references decreases. It is argued that this stabilizing effect is largely a product of uncertainty reduction and normative pressure.

SOCIAL FORCES AND CONSTRAINT IN THE ATTAINMENT OF COMMUNITY STATUS

There exists in the literature on social status a tension between the forces that sort members of a community into differing social strata and the forces that constrain membership in relatively stable status orders. Much has been written about both topics, and yet there has been little empirical research investigating the intersection of these two competing stimuli in shaping the social structure of the community. The objective of this paper is to examine empirically the general proposition that the social forces that create status mobility are the same social forces that, over time, work to establish status stability.

Previous studies of status mobility have been constrained by the use of aggregate evaluations of an actor's status. In other words, an actor's perceived status has usually been assumed to be a stable property among all actors in the community. This approach tends to ignore heterogeneity in status beliefs within the community while also failing to specify empirically the underlying causes that create rigidity within the system. What we would expect to observe in a system of status stability is an increased amount of consistency with respect to the *fundamental inputs* that determine an actor's status. In this study, the generation of an actor's status position is analyzed at the basic level at which status is derived: social approval within the actor-to-actor dyads that are the basis of the community. If status positions do move towards stability, it is at this level of analysis where we would expect to find a convergence of status beliefs.

THE ATTAINMENT OF SOCIAL STATUS

Social status is often defined in terms of an individual's position in a social system. However, in this paper, we will refer to status in the same spirit as Weber (1968), who used the term status in reference to an individual or group's prestige or honor, and Blau (1964), who used the term in reference to actor's claim to respect and compliance in relations with others.

The attainment of social rank and status has been an important topic in the study of groups. Homans (1950) asserted that individuals who generate better sentiment among peers, those whose activities conform closely to the norms of the group, and those who have a wide range of interactions will attain high social ranking. Blau (1964) hypothesized that individuals who make essential contributions to the group, or to its members individually, have a claim to superior status since others in the group are willing to exchange deference in return for the benefits that the higher status individuals have to offer. Thus, status can be thought of as *a product of others' subjective evaluations* of an actor.

These subjective evaluations are often based on observation of an actor's relative position and relationship to others (Frank 1985). Ridgeway and Glasgow Erickson (2000) found that third-party status beliefs spread through the observation of behavior between other individuals. Even the mere presence of a tie between parties can convey status. For

example, in finding that young companies endorsed by prominent exchange partners performed better than otherwise comparable ventures that lacked prominent associates, Stuart, Hoang, and Hybels (1999) hypothesized that much of the benefit of having prominent affiliates stems from the transfer of status that is an inherent by-product of inter-organizational associations. Thus, status can also be thought of as *a product of an actor's position* within its social structure.

Why do actors seek status? Perhaps because, once established, high status can become an asset for its holders (Fombrun 2001). First, high status actors are often given more credit than are low status actors in return for the same amount of effort (Merton 1968) because, as an actor's status increases, the propensity of others to over-estimate the quality of his performance also increases (Sherif 1966). Further, once an actor's role in a domain has been defined, that role becomes a means to claim, bargain for, and gain membership in the larger community of actors (Baker and Faulkner 1991). Therefore, status can be an important factor in determining an actor's pattern of exchange relations. Thye (2000) found that, compared to low status subjects, high status subjects are most often chosen as preferred exchange partners. Finally, when considered with the width of its niche, possessing high status can contribute to an actor's success relative to other actors in the market (Podolny, Stuart, and Hannan 1996).

THE STABILITY OF THE STATUS ORDER

Implicit to the generation of status among actors is the ordering or grouping of social entities along some sort of scale or continuum in which actors are positioned relative to one another in an observable distribution of prestige (Berger, Wagner, and Zelditch 1992). A noteworthy feature of status orders is that, due to structural and institutional forces, these distributions are likely to become characterized by stability.

Structural forces. In the structural approach to markets, structures exist and reproduce themselves in part because the information needed to pursue maximization and efficiency is not available (Leifer and White 1987; White 1981). The only tangible guidance available to actors is information that can be inferred from the patterns and outcomes that emerge from the endogenous positional and social relations among actors. Therefore, it is reproducibility, rather than efficiency, that becomes the relevant issue when defining social structure. Podolny (1993) argued that a market could be thought of as a reproducing status hierarchy in which each producer's status position circumscribes its actions by providing a unique cost/benefit profile for goods produced at a given level of quality. According to Podolny, markets are constrained by processes of status differentiation which Merton (1968) called the "Matthew effect"[1]. According to this

[1] The term is derived from the New Testament Book of Matthew (25:29), which states, "For everyone who has will be given more, and he will have an abundance. Whoever does not have, even what he has will taken from Him."

principle, high status actors obtain greater recognition for performing a given task than do low status actors. As such, high status actors see a greater return on investment for a given level of output than do low status actors and, therefore, accrue a long-term advantage as the benefits of high status increasingly go to those who already have it at the expense of those who do not.

Institutional forces. An actor's status is also subject to institutional pressures. In studying the prestige of firms in a wide variety of industries, Ager and Piskorski (2001) suggest that, over time, a firm's prestige within an industry becomes increasingly more stable as the firm is able to establish a set of reliable expectations of future performance based upon knowledge of it's past behavior. This "taken-for-grantedness" eventually leads to a decoupling between perceived and actual firm quality. As a result, evaluations of prestige ultimately become inert to the extent that changes in an actor's quality are not reflected in changes to an actor's prestige. Berger and Zelditch (1998), claimed that once a prestige order has emerged, others' underlying performance expectations of an actor determine subsequent behaviors (in the giving of resources, for example) that maintain the initial expected performance differences. Since performance expectations and behavior become self-reinforcing entities, the resulting prestige order becomes stable and, over time, an actor's status position becomes increasingly stable.

SOCIAL CUES AS DETERMINANTS OF STATUS

An individual's status in a small group is determined by other group members' expressions of approval (or disapproval) for that individual. However, in a large group or community setting, the number of actors present in the environment may be sufficiently large that each member of the community does not actually know each of the other community members, thus making it difficult for actors to evaluate one another's social status. One of most potent forms of reducing this sort of uncertainty is through social interaction, due to the informational aspect of social contact and the normative forces that ensue (Ross and Nisbett 1991).

Social Cues and Information Transfer. Others are among our best resources for information about the world. Asch (1951) found that independence and yielding in individual judgment are influenced by two factors. First, as the clarity of the stimulus received decreases, the likelihood that an individual will look to the group for clarification increases. Second, as the degree of unanimity within the group increases, so does the likelihood that the individual will yield to the opinion of the majority. At the organizational level, Haunschild (1994) found that firms look to their interlock partners when deciding how much to pay in acquisition premiums—an effect that is stronger under conditions of uncertainty. Davis and Greve (1997) found that firms adopt certain governance changes through a process of board-to-board diffusion. In addition to direct social contact, mimetic behavior is also important. Greve (1996) found that uncertainty

could influence a radio station to adopt a new strategic position if the station merely observed others adopting similar positions.

Social ties can be an especially important source for generating information outside of one's immediate social cluster (Granovetter 1973; Granovetter 1974). The use of others to generate information may be especially important in contexts where information is scarce. Rangan (2000) claimed that, in the presence of spatial dispersion, independent exploration will tend to be viewed as problematic and, therefore, social networks will tend to matter in the search for information. An interesting implication of information transfer through social networks is that, in a community where networks are highly clustered, there is a possibility that information can transfer between clusters and, therefore, between actors in a decidedly non-linear fashion (Milgram 1967; Watts 1999).

Although these studies have not explicitly examined the influence of social cues on social status, the underlying mechanisms of uncertainty reduction and information gathering seem well suited to the study of status generation. Since social status is a product of others' subjective evaluations of an actor, we can think of public expressions regarding a focal actor's status as primary elements of that actor's status (Goode 1978). It follows that, as the number of community members who express an opinion regarding an actor's status increases, information about that actor is more likely to spread throughout the community. As this information becomes more abundant, others will become more likely to be influenced by these opinions:

HYPOTHESIS 1a: As the number of others who give high-status (low-status) references about an actor increases, the likelihood of that actor receiving future high status (low-status) references increases.

The Self-Reinforcing Nature of Status References. Patterns of status references from others should eventually exert a stabilizing force on the focal actor's reputation. As more status-bearing references are received, the more solidified an actor's reputation should become as more information regarding her status position becomes available to the community. Berger and Luckmann (Berger and Luckmann 1966) proposed that within a society, knowledge first becomes institutionalized and then legitimated in a process of social diffusion in which actors come to ascribe cognitive validity to institutional orders—a process by which reality is constructed socially via endogenous social forces. If this is the case, then the receipt of status references should generally lead to a self-reinforcing cycle of endorsements, thus creating stability in an actor's status. Although support for H1a would offer some evidence of a self-reinforcing cycle, analyzing the effects of social cues in both directions is a much stronger test. In other words, if stability is increasing, we should be able to observe a decrease in the amount of contradictory information that enters the community. Accordingly, we should expect that, as the number of others providing high opinions of a focal actor's status increases, the likelihood of that actor receiving low-status references should decrease, or vice versa. Thus, our original hypothesis can be re-stated in a strong test of the stabilizing effects of social cues.

HYPOTHESIS 1b: As the number of others who have given high-status (low-status) references about an actor increases, the likelihood of that actor receiving a low-status (high-status) reference decreases.

The Role of Normative Forces. Hypothesis 1 is based on the proposition that community members rely on information from others in deciding how to assess any particular actor. Besides carrying information, social interactions also carry a great deal of normative pressure. The aggregate judgment of a group creates a strong influence that can push individuals in the group towards an expressed unanimity in opinion (Newcomb 1943). This normative influence is also prominent among sets of organizations which, faced with the same set of environmental conditions, are subject to a process of homogenization (Dimaggio and Powell 1983). In arguing that normative forces influence purposive organizational action, neo-institutionalists have argued that organizations that incorporate societally legitimized elements will maximize their own legitimacy and, therefore, their survival capabilities (Meyer and Rowan 1977), while those who deviate from accepted practices may be penalized for their transgressions (Zuckerman 1999).

For various reasons, the opinions of high status community members are likely to carry more weight than are the opinions of low status community members. First, high status community members are more likely to reflect the ideals of the community (Homans 1950), which makes their opinion valuable to others. Second, low status group members are likely to experience cognitive dissonance if their opinions do not match those of high

status members (Festinger 1957; Heider 1958). Early evidence for these propositions came from Torrance (1955), who found that the suggestions of high status actors are accepted in disproportionate numbers, even if incorrect, and the suggestions of low status actors are rejected in disproportionate numbers, even if correct. Since the opinions of high status actors are valued more highly than are the opinions of low status actors, receiving a positive endorsement from a high status individual is highly beneficial to any individual who is trying to gain social acceptance. Blau (Blau 1964) claimed that "an individual's endeavors to gain social acceptance in a group are furthered most by the approval of highly respected group members, since their approving opinions of him influence the opinions of others and thus have a multiplier effect." Therefore, we should expect the opinions of high status community members to have a disproportionate effect on the opinions of others:

> HYPOTHESIS 2a: As the average status of others who give high status (low-status) references about an actor increases, the likelihood of that actor receiving future high status (low-status) references increases.

Conformity and Status Stability. H2a suggests that members of the community are prone to follow the lead of high status others when deciding where a focal actor fits within a given social sphere, which would suggest the presence of conformity within the community. Thus, we should expect references from high status community members to exert a stabilizing force on the focal actor's reputation. As with H1, a very strong test of stability would be to analyze the effects of normative pressure in both directions. Recent

research has revived the notion that community members are not likely to take actions that threaten their position within the community (Philips and Zuckerman 2001). One consequence of this type of status-protecting behavior is that low and (especially) middle status community members may become reluctant to express opinions that contradict those of higher status community members, since such expressions may cause them to be perceived as having qualities that deviate from those of the highest status community members. Therefore, we can predict that, as the average status of others who claim that a focal actor is of high (low) status increases, the likelihood of others saying that that the focal actor is of low (high) status should decrease.

> HYPOTHESIS 2b: As the average status of others who have given high-status (low-status) references about an actor increases, the likelihood of that actor receiving a low-status (high-status) reference decreases.

THE CONTEXT: ADVOGATO.ORG

These hypotheses will be examined in the context of a community of computer software programmers. Advogato.org is an online community of over 5000 individuals who are involved in developing free, or open-source, software. The goal of the free software community is to preserve the freedom to run, copy, distribute, study, change, and improve software (Stallman 1999). Developers of free software write source-code that is made available publicly so that the code can be distributed to and improved upon by others. Some of the more notable free software projects have evolved around the distribution of the openly evolving Linux operating system, but there are, in fact,

thousands of open-source projects including the popular scripting language PERL, the e-mail server Sendmail, and the web server Apache (Axelrod and Cohen 1999).

The Value of Reputation in Free Software Development. The open source development process has often been described as elitist (Lerner and Tirole 2002). In his influential essay "Homesteading the Noosphere", programmer Eric Raymond claimed that reputation is the capital that drives the open source community. According to Raymond, since the open source community lacks much in the way of formal authority or control, the only way in which a programmer can gain status is through peer recognition:

> "…you become a hacker when *other hackers* call you a hacker. A 'hacker' in this light, is somebody who has shown (by contributing gifts) that he or she has technical ability and understands how the reputation game works. This judgment is mostly one of awareness and acculturation, and can only be delivered by those already well inside the culture." (Raymond 1999)

Having a good reputation can be important for those who want to make key contributions to open source software projects. In a study of contributions to the web-server software Apache, Mockus et al. (Mockus, Fielding, and Herbsled 2000) found that a tight set of "core" developers contributed 83 to 91% of key changes while non-key developers were more likely to contribute to peripheral problems. The reputation that programmers build can be valuable assets for those who work on free software. Linus Torvalds, creator of the Linux operation system, was quoted as saying "The cyberspace 'earnings' I get from

Linux come in the format of having a network of people that know me and trust me, and that I can depend on in return. And that kind of network of trust comes in very hand not only in cyberspace" (Ghosh 1998).

About Advogato. Founded in 1999 by computer scientist Raph Levien, the website Advogato.org serves as a virtual meeting place for developers of free software. Developers join the community voluntarily and create online user accounts that they can use to post publicly viewable diaries, share source code, and participate in public forums related to free software development. In the true spirit of free software, even the source code for the web site is made publicly available. There is no special requirement to create an Advogato account. Anyone interested in participating or observing the activity within the community is welcome to join simply by filling out an online form to create a unique user identity. During the period of this study, the community grew in size at a fairly constant rate to just over 5000 members.

Users can choose the level of information they want to provide to the public about themselves. At a minimum, users provide a username. However, users can (and most do) provide a link to their personal website, an e-mail address, and notes describing themselves and their work on free software projects. This information is all made publicly available on each user's unique Advogato.org personal web page. Once an account is created, users can modify their web page by adding in diary entries. There are no guidelines on the use of diary entries, but they are most often used as a medium for

users to relay thoughts on current projects, editorials on the nature of free software, or any other host of topics that are fresh on the user's mind. See Appendix A for a sample personal page.

(Insert Appendix A about here)

A second major part of the Advogato community is the use of "project" pages, which users can create to describe their participation in unique free software projects. Project pages typically contain information on the name of the project, its web address, notes describing the project, and a list of which Advogato users participate in the project and their roles in the project. Once a user is listed as being a project participant, a line is also added to that user's personal page stating which project the user is involved with and what the user's role is within the project.

For this study, the especially interesting aspect of Advogato is its use of a system of peer certifications whereby any member of the community can provide a publicly displayed evaluation of the status of any other community member. Peer certificates are used as the basis of a tiered status ordering consisting of masters (high status software developer), journeyers (middle status software developer), apprentices (low status software developer), and observers (not verified by others as being a contributor to free software development). The published criteria for defining a user's status revolve around that user's talent and dedication to the free software community, with an emphasis on how influential a user's contributions have been to the development of free software projects.

For a more specific definition of the published guidelines used to define each level, see Appendix B.

(Insert Appendix B about here)

Note that the use of a four-tiered status system allows for an even stronger test of the initial hypotheses, since we are able to include middle-status ratings in the analysis. By the same logic used to generate Hypotheses 1 and 2, we should expect that as the number and average status of others who have given an actor middle-status ratings increases, then the likelihood of that actor receiving future middle-status ratings should increase while the likelihood of that actor receiving explicitly high or low status ratings should decrease.

In order to give a peer certificate, a logged-in Advogato member must first visit another member's individual Advogato page and then provide a certificate by choosing among the four status levels in a pull-down window that appears at the bottom of the receiver's page. Once a certificate is given, it is displayed in two places. First, the certificate is shown on the receiver's page in a list that reflects all certificates received, from whom each certificate was received, and the level of each certificate. Second, the certificate is shown on the sender's page in a list that reflects all certificates sent, to whom each certificate was sent, and at what level they were given. Once a peer certificate is given, it is active indefinitely, unless the provider decides to remove the certificate, which can be done by the same menu-driven process that was used to give the certificate.

Once peer certificates are received, each certificate is weighted by the status of the sender and then combined with all other certificates received in order to calculate a publicly displayed status rank that corresponds with the overall quality of certificates that the member has acquired. The algorithm is designed in a manner that ensures that those members who are successful in attracting high-level peer certifications from high status community members (masters and journeyers) are the most likely to gain high status themselves. Therefore, it is not enough to merely attract high status endorsements—these endorsements must come from the right people. For a more complete description of the metric used to determine a member's Advogato rating, see Appendix C.

(Insert Appendix C about here)

Certificates were designed as a way to confirm whether or not a user is a legitimate member of the free software community. Therefore, each new user is given, by default, the rank of "observer". The only way in which a member can gain a ranking beyond "observer" is to be recognized as such by another user who already holds a rank higher than "observer". In other words, a new "observer" must be certified as being an "apprentice" or higher by at least one user who already holds the rank of "apprentice" or higher. This system of peer certification is designed to protect the integrity of the system by deterring attempts at status mobility by the use of bogus accounts[2].

[2] For instance, it would be difficult for a would-be hacker to generate status by creating a set of bogus accounts in order to manufacture peer certificates. This scheme would work only if the hacker was able to convince a legitimate community member to certify the hacker at a rank of apprentice or above. According to the administrator of the site, due to

Once a member gains the rank of "apprentice" or above, that individual is given the right to contribute to the online forums that are a central part of the Advogato community. These members can use their account to post "articles" on various free software subjects. An article is a text message that is posted on the opening page of the Advogato web site for all community members to read. Those community members who hold the rank of "apprentice" or above are able to respond to the original article if they wish. Thus, a string of lively discussions can ensue, but only amongst those who have been verified as being true "members" of the free software community—simply holding an account as an "observer" does not entitle one to participate in the discussion forums.

METHODS

The Advogato community had over 5000 members at the time the data was collected. As such, it seems highly unlikely that any individual member is familiar with all other community members. Therefore, status attainment should be at least partially dependent upon a social diffusion process in which community members provide publicly expressed evaluations of other community members. In the context of Advogato, peer certificates can be used as evidence that a status evaluation has been given from member j (alter) to member i (ego). Specifically, this study examines the likelihood that alter provides a peer certificate to ego at specific levels within Advogoto's status hierarchy. Those certificates

---

the nature of the algorithm used to calculate overall rank, this type of deception happens very infrequently and with only limited results.

that are made at the level of "master" are clear evidence of high-status approval. Likewise, a peer certificate given at the level of "journeyer" is an indication of middle-status conferral[3] and a peer certificate at the level of "apprentice" or "observer" is indicative of a low-status conferral.

Model and Dependent Variable

The data on individual peer certificates are available in panels that show every certificate given and received by each community member as of time t. Therefore, it is possible to first construct a data set consisting of each member-to-member dyad at time t and then estimate the probability that community member i (ego) receives a peer certificate from community member j (alter) at that time, using the time since the birth of a dyad (the time elapsed since the date that both members existed in the data set) as the time axis. Because the exact form of time dependence for the process of peer certification is not known, peer certifications are modeled using the piecewise constant exponential model, a flexible version of the standard hazard rate model that allows the hazard rate to vary with time in a fairly flexible way and also to vary non-proportionally with covariates. The basic idea of the piecewise exponential model is to split the time axis into discrete time

[3] This is not to say that a middle-status certification means the same qualitatively to different community members. For instance, a peer certification as "journeyer" could be considered as social approval for a community member holding the rank of "apprentice" but disapproval for a member with the rank of "master".

periods and to assume that transition rates are constant within each of these intervals but can change between them. Given time periods $\tau_1,\tau_2,\ldots,\tau_L$, a general form of the model is:

$$r(t)_{ij} = \exp(\acute{\alpha}_l^{(ij)} + \mathbf{A^{(ij)}} * \alpha_l^{(ij)}) \qquad \text{if } \tau_l \leq t \leq \tau_{l+1}$$

where the dependent variable $r(t)_{ij}$ is the transition rate of member i receiving a peer certification from member j at time t, $\acute{\alpha}_l^{(ij)}$ is a vector of constant coefficients associated with the *l*th time period, $\mathbf{A^{(ij)}}$ is a matrix of covariates, and $\alpha_l^{(ij)}$ is an associated vector of coefficients that shows the effects of these covariates in the *l*th time period. This particular model, first proposed by Tuma (Tuma 1980), provides maximum likelihood estimates that allow the effects of covariates (i.e. their associated parameters) to vary across time periods (Blossfeld and Rohwer 1995).

In order to test specific hypotheses, the peer certification data is disaggregated into those certificates that represent high status (certificates given at the level of "master"), middle status (certificates given at the level of "journeyer"), and low status endorsements (certificates given at the level of "apprentice" or "observer"). The disaggregated measures are then used as dependent variables in separate models that estimate the probability that ego receives a specific level of peer certificate from alter at time t_1.

Independent Variables

Hypothesis 1a. In order to test Hypothesis 1, the model includes separate measures for the number of k others ($k \neq j$) who had previously given ego a peer certificate (as of t_0) at the level of master, journeyer, and apprentice, using the observer category as a reference

group. For each model, we predict that the likelihood of ego receiving a certificate at the specific level of the dependent variable increases with the number of others who have previously given ego a certificate at that particular level.

Hypothesis 2a. The average status of the k others ($k \neq j$) who have previously certified ego at each status level was computed by weighting each peer certificate (certificate from an observer=1, apprentice=2, journeyer=3, and master=4) and then dividing this weighted score by the number of others who had given ego a certificate at each level as of t_0. For those individuals who had not received a peer certificate of as of t_0, this variable was set equal to zero. As before, we expect that, as the average status of those who have certified ego at a given level increases, so does the likelihood that alter will also certify ego at that particular level.

Hypothesis 1b and 2b. The stabilizing effects of social cues are tested using the same variables used to test Hypotheses 1a and 2a. However, instead of examining the effects of ego receiving certificates at the same level as the dependent variable, Hypotheses 1b and 2b examine the effects of ego receiving a certificate at a level *other than* the dependent variable. If status stability emerges, we should expect that as the number and average status of those who have previously given ego a certificate at a level other than the focal level increases, the likelihood of ego receiving a certificate at the focal level decreases.

Control Variables

It should be noted here that there are not explicit controls for ego's rank in the model. This is due to collinearity between ego's rank and the variables for the number and status of others who have previously certified ego. The Advogato trust metric is a direct function of these variables and, therefore, most of the variance attributable to ego's Advogato rank will be captured by the number and average status of his certifiers.

Characterstics of alter. There is, however, a control for differences in alter's rank. In all models, indicator variables were included for alter's rank at t_0 with the "observer" category omitted for comparison. Furthermore, it may be that some community member's are just more prone to giving certificates than are others. Therefore, the models include a term for the total number of certificates that had been given by alter as of t_0. Similarly, community members may differ in their inclination for giving certificates to high or low status others. For example, an individual who certifies mainly high-status others would have a lower baseline likelihood of giving certificates to members with low status. Therefore, the models include a variable that controls for the average status of those whom alter had previously cited as of t_0.

Ego's tenure. Although we expect the power of social cues to be strong, it seems likely that the strength of these effects will decrease with time. If ego's status position does become stable, there may be a temporal effect created by the institutionalization of his status position. Simply put, once ego's social position becomes institutionalized, there

may come a point at which the issue of his status ceases to be a focal point for discussion. Therefore, the models include a control for ego's tenure, as measured by the length of time, in days, since ego created his Advogato account and, therefore, became an official member of the Advogato community.

Alter's tenure. There is also a variable for alter's tenure, since there may be negative tenure dependence in alter's motivations for giving a status reference. Actors within a community can use ties to others as a way to establish their own social identity (Podolny 2001). In the early stages of community membership, alter may be more likely to use peer certifications as a signal to other community members, in essence broadcasting what it is that she believes her own social position to be. However, once her social position becomes established, there is less of a signaling benefit in providing certifications to others, thus leading to a decline in her motivation to give status references.

Reciprocity. A cursory glance at individual web pages reveals that there are a good number of reciprocal certifications in this community. In other words, there are numerous instances in which ego and alter cite each other. It seems highly feasible that, if ego has already certified alter, the norms of reciprocity and exchange may come to dominate the relationship, thus affecting the impact of external social cues. Therefore, the models include a variable indicating whether alter had been certified by ego as of t_0. It also seems reasonable that a recent cite from ego should have a more powerful effect than one given long ago, with the likelihood of alter reciprocating being the strongest

immediately following the receipt of a certificate and decreasing in a non-linear fashion. Therefore, the models also control for the logged number of days since ego first gave a peer certificate to alter.

Attention effects. In exploratory models, it was found that, conditional upon ego receiving a peer certificate of any level, there was an "attention" effect on the probability of ego receiving further peer certificates of any type. In other words, because some community members are simply more visible than others, a peer certificate received by a highly visible individual is empirically associated with a higher likelihood of receiving future peer certificates of any type, regardless of the level at which the focal certificate was given. Therefore, in an attempt to control for this attention effect, it was necessary to include a fixed effect for the total number of peer certificates received by each ego. Because a simple count variable would be collinear with the disaggregated by-type reference counts that we use to test Hypothesis 1, the models include multiple categorical variables that control for the number of previous cites that ego had received as of t_0, using the modal category of 1 as the reference group. Since the number of previously received references had a maximum value of 527, including a fixed effect for each value would have literally generated several hundred new variables, thus making the models impossible to estimate given the capacity limitations of the computer software technology currently available. Therefore, this variable was aggregated into a smaller number of discrete categories. Most community members had only a small number of peer certificates at t_0, so the first 19 categories (which included only 1 value each, i.e. the

number of references at t_0=1, =2, =3, etc.) accounted for 87 percent of the observations. The remaining values (number of references at t_0=20 or more) were aggregated into groups that accounted for roughly 1 percent of the data in each group. Since the total number of references is a linear function of the number of master, journeyer, apprentice, and observer references, one of these variables had to be omitted from the models and used as a reference group. In order to produce the most easily interpretable contrasts, the variable for the number of "observer" certificates received as of t_0 was omitted.

Creating Analyzable Data. Data were obtained from the original .xml data files that were used to generate the Advogato.org website. The creator and current administrator of the Advogato site provided the data by supplying 13 discrete panels created at roughly one-month intervals from July 2000 through August 2001. In each panel, there existed a separate file for each individual community member that mirrored the data available on each community member's personal Advogato web page on the date of the panel. In addition, there was an assortment of information not available over the web, such as the exact date when that user created his or her Advogato account. Due to the comprehensive nature of the panels, the size of each panel was enormous. In order to manage the large amount of data, a custom-written java script was used to extract the pertinent data from the .xml code and then insert the data into a SQL database, from which it was possible to use standard query language to combine information from the discrete panels into a single ascii data file.

The resulting ascii data file was then used as the input source for creating the analyzable data set. In order to analyze the data at the level of each member-to-member dyad, panel data was re-constructed by creating unique observations for each i-to-j dyad combination that was possible as of time t. In other words, if both ego and alter were alive on or before time t, there existed the possibility of peer certification between the two and, therefore, a unique i-to-j observation was inserted into the data set. Thus, as the number of community members increased, the number of possible dyads increased exponentially. Because the estimation of our models required lagged values for the independent variables, the first observation of each dyad was not informative since there was no lagged information available for at least one of the two members. Therefore, the first observation from each dyad was excluded from the analysis.

Unobserved heterogeneity. In the original data set, there were several million observations created by community members who were inactive in peer certification. During the period of the study, approximately 40 percent of the population neither gave nor received any peer certificate during their tenure in the community. One concern in leaving these individuals in the risk set is that they may have chosen at their own discretion to be inactive members of the community and, therefore, could be a source of unobserved heterogeneity. For instance, alters who have voluntarily excluded themselves from participation might fail to give a certification to ego, but not because they are immune to social influence. Rather, they may have simply opted not to participate in the Advogato certification system.

A second concern is that, because there are no restrictions or formal filtering processes in the creation of new user accounts, some of the individuals who never receive a certificate are simply not legitimate members of the free software community and, therefore, are not truly members of the population at risk for receiving a peer certificate. If this is the case, leaving these non-participant individuals in the risk set could lead to bias in our models since we would be comparing legitimate "at-risk" community members with individuals who are never really at risk of receiving a peer certificate. For these reasons, those individuals who neither gave nor received a peer certificate during the period under observation were excluded from the study—in essence making this a study of social influence only among actively participating community members[4].

For each dyad, a starting time was given as the date upon which the latest member (ego or alter) joined the community. Because no Advogato accounts were removed during the period of the study, all dyads were right-censored as of the last panel. For each dyad,

[4] A possible drawback of this strategy is that it leaves open the possibility of differential retention. In other words, since non-participants are excluded from the analysis, there is a possibility that what we observe as social influence is really just the effect of having those whose opinions differ from majority choosing not to participate in peer certification. Qualitatively, we have no reason to believe that this is the case. Moreover, given the lack of data on these individuals, this would be extremely difficult to refute empirically.

covariates were updated and a new observation was added in each subsequent panel, which resulted in an initial data set containing 54 million observations. Removing observations in which there was no new or changing data from t_0 to t_1 reduced the total number of observations to 43 million. Unfortunately, due to capacity limitations in computing technology available at the time of the analysis, this data set was much too large to analyze in its entirety, especially when accounting for the spell-splitting that occurs in the piecewise exponential model. Therefore, it was decided that a pseudo-random sample would be drawn. After drawing samples of different sizes, it was found that a sample of 33% of the original data (sampled by ego's id number) would be small enough to analyze without exceeding the limitations of the available computer hardware and statistical software. Comparing descriptive statistics between the full data set and multiple samples of 33% revealed only minute variance across samples in the mean and standard deviations of all variables, suggesting that the sample size chosen is a robust representation of the full data set. Within the sample, there were nearly 14 million observations representing roughly 6 million unique dyads.

Obtaining Baseline Estimates of Time Dependence.

It seems plausible that there should be negative duration dependence within each actor-to-actor dyad. If alter enters the community already holding knowledge about ego, then the likelihood of ego receiving a status reference from alter may be very high early in the life of the dyad. If, on the other hand, there is uncertainty in the relationship, alter might choose to wait for some amount of time in order to see what others are saying about ego.

If alter does not give an endorsement during this initial information-gathering period, then it seems unlikely that he will ever give a reference to ego unless new information about ego should come his way.

Indeed, a cursory examination of the Nelson-Aalen cumulative hazard estimates revealed that the hazard rate of ego receiving a peer certificate from alter at each level was greatest in the earliest days of a dyad's history and then declined slowly for most of the remaining life of the dyad. Based on visual inspection of the cumulative hazard rates and upon the belief that the hazard of receiving a certificate declines over time, the exponential hazard model was broken into pieces at the intervals of 50, 100, 200, and 400 days resulting in five time periods. An analysis of the baseline model with no covariates (not shown) confirmed that the hazard of receiving a peer certificate declined gradually as a dyad's age increased. The results of baseline models for each of the disaggregated dependent variables also showed this same form. Figure 1 shows the graph of the cumulative hazard rate for ego receiving a peer certificate.

(Insert Figure 1 about here)

RESULTS

Descriptive statistics are provided in Table 1 and results of the multivariate analyses are reported in Table 2. The control variables included in this study were generally very strong and robust predictors of alter giving a peer certificate to ego. Advogato ranking was an important predictor of alter's behavior as was alter's tenure and previous tendencies in giving peer certificates. Although this study does not speculate the reasons

why, it seems as though some community members are just more prone to giving peer certificates than are others. As the number of cites given by alter to k others ($k \neq j$) as of t_0 increased, so did the likelihood that alter would give a certificate to ego at t_1. There was also a very strong effect for the average status of those whom alter had previously certified—an effect that was especially prominent in the models for master and journeyer-level certificates. Of all the control variables, reciprocity had the most significant effect. The presence of a certificate from ego to alter had a massive positive effect on the likelihood of alter giving a certificate back to ego, although a return certificate was far from automatic. It is interesting to note that the effects of reciprocity were the strongest in the models of middle-level status certificates.

(Insert Table 1 about here)

Across all three sets of models, Hypothesis 1a received very strong support. In all models, an increase in the number of others who had previously given ego a specific level (master, journeyer, or apprentice) significantly increased the likelihood that alter would give ego a certificate of that same type ($p<0.001$). Thus, it seems likely that, within this community, alter uses information from others in forming an opinion of ego, lending support to the more general sociological proposition that status beliefs are spread via social diffusion.

It also appears as though alter weights the information from others according to their status. As predicted by Hypothesis 2a, in all models, the average status of others certifying ego at a given level had a highly significant positive effect on the likelihood

that alter would give a certificate of that same type. Thus, in this context there also appears to be a strong normative influence in the social diffusion of status beliefs.

Hypotheses 1b and 2b also received strong support in all models. Not only did the number and average status of certifiers at a given level have a positive influence on the likelihood of alter providing a certificate at that same level, but the number and average status of others giving certificates at a level *other than* the focal level significantly decreased the likelihood that alter would give a certificate at the focal level. In sum, as information regarding ego's status diffuses through the community, the likelihood of ego receiving a peer certificates that runs counter to the established pattern of peer certification decreases, lending support to the proposition that social approval and disapproval from others exerts a stabilizing force on the focal actor's status. Put simply, high status certificates tend to inhibit downward mobility, low status certificates tend to inhibit upward mobility, and middle status certificates tend to inhibit both upward *and* downward mobility.

(Insert Table 2 about here)

This pattern is consistent with sociological theories regarding the interaction between status and social diffusion. For instance, Goode (Goode 1978) proposed that the higher the rank of a person who criticizes others, the greater the likelihood that others will agree with the criticism. We find direct support for Goode's proposition in the models of low-status citations and indirect support in the models of high status certificates. Not only did disapproval (a low-status certificate) from high status others exert a positive influence on

the likelihood of alter expressing disapproval, but high-status disapproval also exerted negative pressure on the likelihood of alter giving approval (a high-status certificate). Table 3 displays a summary of the independent variables and the direction of their estimated coefficients.

(Insert Table 3 about here)

It should also be noted that, ceteris paribus, ego's tenure had a very strong negative effect on the likelihood of receiving a certificate of any type ($p<0.001$). While the strong negative tenure dependence does not provide direct evidence of status stability, it does suggest that others become less likely to make overt statements regarding a focal actor's status as his tenure increases, which leaves open the possibility that, as an actor's status position becomes increasingly institutionalized or "taken-for-granted", other community members feel less compelled to talk about his status.

Moreover, alter's tenure also has a strong negative effect ($p<0.001$). The longer alter is in the community, the less likely it becomes that she will certify ego. One feasible interpretation is that alter becomes less likely to give overt approval or disapproval as her own status becomes increasingly well established. Since alter is most likely to provide certifications when her own tenure is short, perhaps it is at this point when the act of giving a certification is most beneficial. In other words, alter can use certifications as a tool for publicizing her own social position. As such, in the early stages of community membership, alter may be more likely to use peer certificates as a signal to other community members to broadcast what it is that she believes her own social position to

be. However, once her social position becomes established, there is less of a signaling benefit in providing subsequent certifications.

DISCUSSION

One weakness with this data set is that we do not know if ego and alter actually know each other. The presence of a relationship between ego and alter would likely have a dampening effect on the social diffusion process. It seems likely that if ego and alter know each other, then the amount of status uncertainty between them should be comparatively lower than the amount of status uncertainty between two strangers. As a result, the influence of outsiders in the process of reputation formation could be mitigated to a certain extent.

One possible solution would be to integrate data regarding the projects that ego and alter work on. If both work on the same project, then there is a chance that they know each other. However, given the nature of open-source software development, the probability that they know each other based upon project contributions alone is still not very high. Open source projects rely on voluntary contributions from users who work and live in a dispersed, international environment. Furthermore, since many projects are only loosely organized, contributors can contribute freely to projects without knowing or communicating with other contributors, often to the point of not even knowing who else has contributed to the project.

So perhaps there is a better proxy for acquaintanceship. One possible measure of familiarity is reciprocity. It is interesting to note that new memberships in Advogato often occur in social clusters. One of Advogato's most active members was quoted as saying:

> "A bunch of people from some community (e.g., Bay Area Linux user groups, employees of a particular company, developers on a particular project) will get involved, and all enjoy reading one another's diaries all the time. That keeps them coming back…People have certainly come and gone in waves." (Chalmers 2000)

The idea that community membership grows in identifiable social clusters was also confirmed in personal discussions with the Advogato.org founder, who noted that new accounts are often formed in batches of pre-existing social cliques. This clustering would be consistent with the negative dyadic tenure dependence that occurs in the baseline models. In other words, if new memberships are often created in pre-defined social clusters, these members arrive in the community with pre-existing knowledge about the status of particular others, which should lead to shorter wait times for peer certification within these dyads.

If this is the case, then perhaps the most reliable measure of a pre-existing relationship is the presence of a reciprocal certification. Reciprocity had a huge influence on the hazard rate for peer certification. For example, in the model of journeyer-level certificates, the presence of a certificate from ego to alter made alter nearly 45 times more likely to

certify ego (44.87=exp(3.8038), $p<0.001$). The strength of this effect suggests that those members who exchange certifications may be more likely to know each other than those who do not. Even though it is far from perfect as an identifier of pre-existing relationships, the presence of a reciprocal certification is probably the best proxy available in the data at this time.

Lack of quality measures. Another omission in this data is the lack of concrete measures to control for the actual quality of ego's work or ability. While this study suggests that ego's reputation is socially constructed, there should be some amount of variance that is explained by underlying differences in the actual quality of ego's contribution to the free software community. In theory, all Advogato certifications should be made according to the criteria listed in Appendix B. Unfortunately, we don't know how often members actually refer to these criteria when making peer certifications. As such, there is no real objective measure of the quality of ego's contributions available in the data.

However, there is good reason to believe that ego's quality may be partially accounted for through his affiliation network. Podolny (Podolny 2001) has suggested that a tie between two actors is an important informational cue that others rely upon to make inferences about the underlying quality of one or both of the market actors. In Podolny's metaphor, a tie between actors serves not only as a conduit for information flow between actors (like a pipe), but also as an apparatus that produces differentiation between actors by affecting others' perception of an actor's relative quality (like a prism). Thus, in an

environment in which there exists at least some amount of uncertainty, community members will rely on the status of those with whom an actor has ties as a proxy for the underlying quality of the focal actor's services. Podolny's "market as a prism" metaphor is somewhat related to the economic theory of market signaling (Spence 1974), in which market sellers send buyers signals that convey information about a product's quality in order to deal with problems of information asymmetry. Within Advogato, the presence of a peer certificate from alter to ego establishes a tie that could be interpreted as a signal of ego's underlying quality. In other words, members of the community could be using the status of those who certify ego as a signal of ego's true quality.

One important aspect of signaling is that the value of the signal depends on the ease with which the signal for high quality can be obtained by lower-quality actors (Milgrom and Roberts 1992). Therefore, in order for certificates to have positive value, high status certifications from high status others should be difficult for low-quality actors to obtain. Although this analysis does not provide definitive support for the presence of signaling via alter's status, there are patterns within the data that support this conclusion. For instance, in models of reciprocated master-level peer certificates (not shown), it is those community members of high status who are, by far, the least likely to give out reciprocal high-status certificates. Thus, although it is not a perfect substitute, peer certificates fit the general conditions of market signaling, leaving open the possibility that the status of others who certify ego can serve as a proxy for ego's unobserved quality. Even though

this does not provide a perfect control for quality, it will have to suffice until such time as a more objective measure can be included in the analysis.

In the end, however, Advogato remains a pure reputation game. This is by design and implementation within the trust metric. The trust metric is basically a tool that measures not the objective quality of an individual's work or ability, but others' subjective beliefs about the quality of her work and ability. Thus, community members attain status in the Advogato realm community not necessarily by *being* good programmers, but by having others *say* that they are good programmers. Of course, this is not to say that there isn't a correlation between the two. However, there is no way of specifying what or how strong that correlation is.

Future Research. Besides the addition of quality measures, there are other analyses that will be useful in future studies. As mentioned previously, one area in which this data might prove to be particularly fruitful is the study of reciprocal certificates. The results of preliminary models (not shown) that focus on reciprocal certificates suggest that role of social influence remains significantly strong even in the presence of reciprocity. However, there remain untested several interesting interactions between ego's rank, alter's rank, and the level of certificate given/received that could prove to be important determinants of the likelihood of alter returning ego's certification. It might also be useful to examine whether reciprocal certifications are exchanged under the norms of direct (or restricted) exchange, or rather under the norms of generalized exchange in

which alter may choose to provide reciprocity indirectly (Bearman 1997; Takahashi 2000). In order to test this matter, a more thorough examination of the social structure of the larger community would need to be established. One possible route for determining how community members are linked to one another outside of the realm of peer certificates is the use of the project data. There exists data on approximately 1100 open-source software projects. Unfortunately, the reliability and validity of the project data is questionable, since there appears to be much variance in the amount of care and effort put into the creation and maintenance of the project pages. However, these web pages contain at least some information on the staffing of each project. This information might be merged onto the peer certificate data to provide a more thorough picture of how community members are clustered within the social space of actual free software projects and whether this affects the pattern of intra and inter-cluster certification behavior.

IMPLICATIONS

In sum, the results of this study present a vivid example of the social forces at work in establishing both status mobility and status stability. In the process of reputation formation, others will tend to evaluate a focal actor's reputation according to socially driven cues provided by others who have previously expressed social approval or disapproval. These factors are important determinants of the ability of an actor to gain the overt expressions of approval necessary to generate status mobility. Thus, it is somewhat ironic that the results imply that these same social forces eventually work to produce status stability. It was found that, as the number and average status of others

who gave high (low) status ratings of an actor increased, the likelihood of that actor receiving a future high (low) status rating also increased. At the same time, the likelihood of that same actor receiving a low (high) status rating significantly decreased. It has been asserted here that this effect is largely a product of uncertainty reduction and normative pressure. As more people certify an actor's status, the amount of information available about that actor increases, thereby reducing uncertainty regarding that actor's status. Furthermore, as that actor draws the attention of high status others, the pressure of community members to conform to the opinions of those high status others increases. Thus, as the number of references (and the average status of those making the references) regarding a focal actor's status increases, the likelihood of that actor being recognized as holding a status that differs from these references decreases.

Moreover, as time progresses, it becomes less likely that other actors will generate overt statements regarding ego's status. There is a very strong tenure effect that reduces the odds of an actor receiving public rating of any type—high status, low status, or indifferent. This finding is congruent with the proposition that, over time, an actor's status becomes increasingly institutionalized or "taken-for-granted", and, as such, other community members seem less prone to making statements about that actor's social position. One very practical implication of these findings is that, if an actor does not begin to establish high status quickly, her odds of doing so may deteriorate with time. The results of this study suggest that speed matters in the establishment of status. Because an actor's position in a status order becomes stable over time, members of a

community who wish to establish high status should do so quickly or they may run the risk of being cast into an inert low status social position.

CONTRIBUTIONS

Although there has been a fair amount of conjecture regarding the circumstances that lead to the generation and stability of status orders, most of the previous work in this area has typically studied the role of stabilizing forces by implicitly assuming that a status order has evolved endogenously, but without specifying empirically the underlying causes that create rigidity within the system. In this study, we observe a complex real-world social system in which status stability is created by an increased amount of consistency with respect to the *fundamental inputs* that determine an actor's status. We find status rigidity being created by the same basic social cues that enable status mobility. Previous studies of status stability have been constrained by the use of aggregate evaluations of an actor's status. That is, an actor's status has usually been assumed to be a stable property among all other actors in the population. Unfortunately, the use of aggregate status measures ignores the presence of heterogeneity in status beliefs within the population. In this study, the generation of status has been examined at the basic level at which status is derived: social references within the actor-to-actor dyads that form the basis of the community. An actor's social status is an aggregated derivative of social approval at this level. Therefore, in a community where status positions become stable, it is at this level of analysis where we find a convergence of status beliefs.

This study also adds to the literature by empirically investigating the role of stabilizing forces in a very large-scale setting. There has been little attention given to the examination of these processes beyond small groups. Implicit to the generation of status orders is the grouping of social entities along some sort of scale or continuum, a task that is made inherently easier with smaller numbers. This study examines the processes of status ordering in a setting in which the number of actors is sufficiently large that actors are unlikely to have information about the majority of other actors. In such a setting, the pattern of emerging social relations is uncertain and potentially difficult to predict (Kollock 1994). Perhaps status stability is a reaction to this uncertainty. In the socially driven status generation process that developed in this large community setting, high-status endorsements tended to inhibit downward mobility while low-status endorsements tended to inhibit upward mobility, thus contributing stability and predictability to the evolving status order.

FIGURE 1. Nelson-Aalen Cumulative Hazard Estimates of the Likelihood of Ego Receiving a Peer Certificate from Alter.

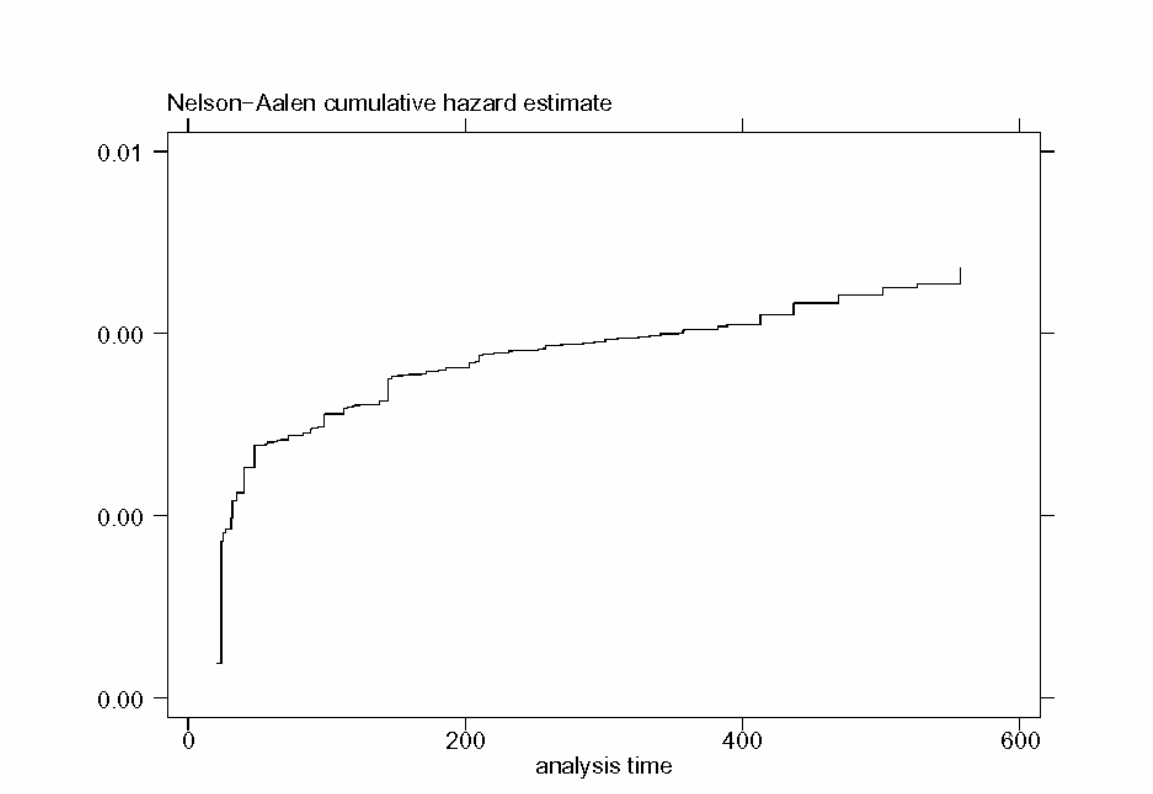

TABLE 1. Univariate Statistics for the Analysis of Advogato Peer Certifications

Variable	Mean	Std. Dev	Min	Max
Variables describing j (alter)				
j's rank at t_0 = observer	0.4033	0.4906	0	1
j's rank at t_0 = apprentice	0.1354	0.3422	0	1
j's rank at t_0 = journeyer	0.3439	0.4750	0	1
j's rank at t_0 = master	0.1173	0.3218	0	1
j's tenure (in days)	309.5300	145.6185	21	667
Number of references given by j at t_0	9.9811	19.1541	0	280
Average status of others rated by j	2.4063	1.4471	0	4
j received certificate from i at t_0	0.0042	0.0648	0	1
time since j received certificate from i (in days)	0.8400	14.8604	0	413
Variables describing i (ego)				
Number of master certs received as of t_0	3.8632	18.6125	0	496
Number of journeyer certs received as of t_0	5.4052	8.9628	0	122
Number of apprentice certs received as of t_0	2.1079	3.6336	0	46
Number of observer certs received as of t_0	0.7822	2.2968	0	47
Average rank of others giving master certs	0.7249	1.2607	0	4
Average rank of others giving journeyer certs	1.7794	1.3894	0	4
Average rank of others giving apprentice certs	1.5332	1.4363	0	4
Average rank of others giving observer certs	1.0448	1.4583	0	4
i's tenure (in days)	315.2918	145.9455	21	667
Dependent Variable				
i received master certificate from j	0.0014	0.0371	0	1
i received journeyer certificate from j	0.0020	0.0442	0	1
i received apprentice certificate from j	0.0008	0.0274	0	1
i received observer certificate from j	0.0002	0.0146	0	1

# of community members = 2836
# of dyads = 5930442
# of observations = 13800773

TABLE 2. Maximum Likelihood Estimates for Models of Ego Receiving a Peer Certificate at t_1

Variable	Master-Level Coef.	Std. Err	P>|z|	Journeyer-Level Coef.	Std. Err	P>|z|
tp1 <= 50 days	-14.5354	0.1899	0.000	-11.8760	0.0856	0.000
50 < tp2 <= 100 days	-14.5420	0.1936	0.000	-12.0017	0.0913	0.000
100 < tp3 <= 200 days	-13.9913	0.1971	0.000	-11.4210	0.0954	0.000
200 < tp4 <= 400 days	-13.1322	0.2047	0.000	-10.3789	0.1049	0.000
400 days < tp5	-12.0941	0.2483	0.000	-8.9126	0.1481	0.000
Variables describing j (alter)						
j's rank at t_0 = apprentice	0.1022	0.0510	0.045	0.2751	0.0445	0.000
j's rank at t_0 = journeyer	0.5376	0.0364	0.000	0.8884	0.0333	0.000
j's rank at t_0 = master	0.5964	0.0443	0.000	0.4882	0.0412	0.000
j's tenure (in days)	-0.0097	0.0002	0.000	-0.0086	0.0001	0.000
# of references given by j at t_0	0.0168	0.0003	0.000	0.0178	0.0002	0.000
Avg status of others rated by j	0.8292	0.0204	0.000	0.4747	0.0132	0.000
j received certificate from i at t_0	3.2441	0.2146	0.000	3.8038	0.1392	0.000
ln time since j received certificate from i	-0.1257	0.0471	0.008	-0.0278	0.0302	0.357
Variables describing i (ego)						
# of master certs received as of t_0	0.0067	0.0002	0.000	-0.0075	0.0022	0.001
# of journeyer certs received as of t_0	-0.0274	0.0014	0.000	0.0518	0.0022	0.000
# of apprentice certs received as of t_0	-0.1569	0.0107	0.000	-0.0450	0.0034	0.000
Avg rank of others giving master certs	0.2941	0.0153	0.000	-0.1098	0.0092	0.000
Avg rank of others giving journeyer certs	-0.1983	0.0122	0.000	0.0435	0.0124	0.000
Avg rank of others giving apprentice certs	-0.0334	0.0109	0.002	-0.1018	0.0085	0.000
Avg rank of others giving observer certs	-0.1728	0.0116	0.000	-0.0464	0.0078	0.000
i's tenure (in days)	-0.0057	0.0001	0.000	-0.0068	0.0001	0.000
# of Dyads	5930442			5930442		
# of Events	6482			9628		
# of Observations	20260243			20253541		
Log Likelihood	-29413.91			-48836.40		
LR Chi-Square	44688.55	50 d.f.		48709.16	50 d.f.	
Prob > Chi-Square	0.000			0.000		

* indicator variables for ego's total number of references not reported

TABLE 2 (cont.) Maximum Likelihood Estimates for Models of Ego Receiving a Peer Certificate

Variable	Apprentice-Level Coef.	Std. Err	P>|z|	Observer-Level Coef.	Std. Err	P>|z|
tp1 <= 50 days	-11.2747	0.0880	0.000	-12.7604	0.1750	0.000
50 < tp2 <= 100 days	-11.5610	0.1003	0.000	-12.6656	0.1952	0.000
100 < tp3 <= 200 days	-11.0775	0.1091	0.000	-12.2643	0.2127	0.000
200 < tp4 <= 400 days	-10.1437	0.1283	0.000	-11.3493	0.2470	0.000
400 days < tp5	-8.5357	0.2200	0.000	-10.2318	0.4073	0.000
Variables describing j (alter)						
j's rank at t_0 = apprentice	0.5278	0.0585	0.000	-0.0278	0.1281	0.828
j's rank at t_0 = journeyer	0.8088	0.0474	0.000	0.6932	0.0942	0.000
j's rank at t_0 = master	0.4122	0.0631	0.000	0.9238	0.1108	0.000
j's tenure (in days)	-0.0062	0.0002	0.000	-0.0054	0.0004	0.000
# of references given by j at t_0	0.0180	0.0002	0.000	0.0218	0.0004	0.000
Avg status of others rated by j	0.1272	0.0150	0.000	0.1704	0.0306	0.000
j received certificate from i at t_0	4.3339	0.2208	0.000	1.9654	0.5686	0.001
ln time since j received certificate from i	-0.0329	0.0487	0.499	0.0616	0.1200	0.608
Variables describing i (ego)						
# of master certs received as of t_0	-0.0015	0.0040	0.710	-0.0098	0.0024	0.000
# of journeyer certs received as of t_0	-0.0073	0.0050	0.143	-0.0471	0.0043	0.000
# of apprentice certs received as of t_0	0.1140	0.0057	0.000	-0.0690	0.0083	0.000
Avg rank of others giving master certs	-0.2547	0.0274	0.000	-0.2992	0.0353	0.000
Avg rank of others giving journeyer certs	-0.3059	0.0163	0.000	-0.2600	0.0285	0.000
Avg rank of others giving apprentice certs	0.2857	0.0165	0.000	0.0659	0.0231	0.004
Avg rank of others giving observer certs	-0.0580	0.0133	0.000	0.3859	0.0222	0.000
i's tenure (in days)	-0.0073	0.0002	0.000	-0.0089	0.0004	0.000
# of Dyads	5930442			5930442		
# of Events	4165			1180		
# of Observations	20268164			20274254		
Log Likelihood	-24891.00			-8699.19		
LR Chi-Square	20953.23	50 d.f.		5426.89	50 d.f.	
Prob > Chi-Square	0.000			0.000		

* indicator variables for ego's total number of references not reported

TABLE 3. Summary of MLE Estimates for Independent Variables and Observed Direction of Coefficients

		Observed Direction of Coefficient by Type of Certificate Modeled			
Hypothesis	Independent Variable	Master	Journeyer	Apprentice	Observer
1a & 1b	Number of others who have previously given ego a peer certificate:				
	Master Certifications	+	-	-	-
	Journeyer Certifications	-	+	-	-
	Apprentice Certifications	-	-	+	-
2a & 2b	Average status of others who have previously given ego a peer certificate:				
	Master Certifications	+	-	-	-
	Journeyer Certifications	-	+	-	-
	Apprentice Certifications	-	-	+	+
	Observer Certifications	-	-	-	+

APPENDIX A. Sample Personal Page from Advogato.org. (Source: http://www.advogato.org/person/zhaoway/)

Personal info for zhaoway

This person is currently certified at Journeyer level.

Name: zhaoway

Homepage: http://www.zhaoway.com/

This person is:

- a Developer on project Debian.
- a Documenter on project Evangelism.

Recent diary entries for zhaoway:

14 Apr 2002 »

Gradually I'm having a different attitude towards correctness proving. But I need to read more. And I need to code more. I heard Erlang is a good real time language and is functional. But I will be far away from mastering it that I would be able to put it into work. Generally, though I feel functional programming is difficult, Clean makes me feel that it is interesting, and could be easier later on, at least I hope so. So I'd like to do more in that. And in the pure way.

Uh, the Object I/O library of Clean, which is difficult. I think the difficulty is mainly because of the type system and the pure nature of Clean. Whatever, it is difficult, albeit the *uniqueness* discipline is pretty simple and straight. I spent the early morning solving a small problem of local state passing, and finally did it. Hee hee.

From Scheme I learnt *continuation* which is by far the most *cool* thing I learnt from programming languages, but I have to say that impure functional languages gives little in helping learning the functional style.

I'm more or less lost in knowing what the hell the benefits of macros are. That means programming is a difficult topic, I mean, only after I do a project which materialize the power of macros, can I really understand the power of macros. My beliefs in macros would be just lies 'till that day comes.

Lone Wolf

I read that Hermann Weyl being called the *lone wolf* among Hilbert's students. It's amazing that an *egg head* mathematician being called *lone wolf*. Previously I thought only characters in Hemmingway or Jack London's books or from Schwazeneiger's movies will ever be called *lone wolf*. Funny. You'd love mathematicians.

Elephant

Okay, I think I could name my warm-up project in the Clean language now. The name is *elephant*. It's a program playing Chinese Chess a la Xiangqi through TCP. Only 10k characters have been written so far. But I did draw all of the 14 pretty pictures of the chessmen using only the mouse and successfully loaded them through the API jungle of Clean's Object I/O library. I am proud of it. Hehehe.

65 older entries...

This person has certified others as follows:

- zhaoway certified foka as Journeyer
- zhaoway certified spacehunt as Journeyer

- zhaoway certified perlamer as Journeyer
- zhaoway certified wichert as Master
- zhaoway certified tausq as Journeyer
- zhaoway certified evo as Journeyer
- zhaoway certified andersee as Master
- zhaoway certified claviola as Journeyer
- zhaoway certified sye as Apprentice
- zhaoway certified bratsche as Journeyer
- zhaoway certified atai as Journeyer
- zhaoway certified ianmacd as Journeyer
- zhaoway certified fxn as Journeyer
- zhaoway certified thom as Journeyer
- zhaoway certified forrest as Journeyer
- zhaoway certified crhodes as Journeyer

Others have certified this person as follows:

- foka certified zhaoway as Apprentice
- spacehunt certified zhaoway as Journeyer
- perlamer certified zhaoway as Journeyer
- jrf certified zhaoway as Journeyer
- fxn certified zhaoway as Journeyer
- sye certified zhaoway as Journeyer
- evo certified zhaoway as Journeyer
- Joy certified zhaoway as Apprentice
- chalst certified zhaoway as Journeyer
- forrest certified zhaoway as Journeyer

Certify zhaoway as: Observer Certify

See the Certification overview for more information.

[Home | Articles | Account | People | Projects]

APPENDIX B. Advogato Certification System (Source: http://www.advogato.org/certs.html)

Master

A Master is the principal author or hard-working co-author of an "important" free software project, i.e. one that many people depend on, or one that stands out in quality. A Master has command of the tools and is an excellent programmer. Generally, a Master works equivalent to full time (or more) on free software. Ideally, a Master writes clearly about the work and its broader context, and serves as a mentor to others in the free software community.

Journeyer

Journeyers are the people who make free software happen. A journeyer contributes significantly to an important free software project, or is the author of a useful or technically innovative project. A Journeyer is generally a competent programmer, but significant contributions of documentation, artwork, or other non-code goodies counts too. Ideally, a Journeyer works with others in the free software community to polish and refine the library of free software. While not necessarily the equivalent of full time, a Journeyer spends a significant amount of time on free software.

Apprentice

An apprentice is someone who has contributed in some way to a free software project, but is still striving to acquire the skills and standing in the community to make more significant contributions. Ideally, the Apprentice is in touch with either an individual mentor or a community that helps to gain these skills. An Apprentice spends a significant amount of time learning the craft of software development, whether by hands-on practice, academic study, or careful observation.

APPENDIX C. Advogato Trust Metric (Source: http://www.advogato.org/trust-metric.html)

Advogato's trust metric

The basic trust metric evaluates a set of peer certificates, resulting in a set of accounts accepted. These certificates are represented as a graph, with each account as a node, and each certificate as a directed edge. The goal of the trust metric is to accept as many valid accounts as possible, while also reducing the impact of attackers.

Advogato performs certification to three different levels: Apprentice, Journeyer, and Master. This is actually done by running the basic trust metric three times, using the "level" value in the certificate as a threshold. Thus, certification of Apprentices is computed using all certificates, while Master is computed using Master certificates only. The computation of the trust metric is performed relative to a "seed" of trusted accounts.

The core of the trust metric is a network flow operation. Informally, if there is a rich web of interconnections, flow reaches almost all the nodes. However, only a few accounts would be accepted from a large nest of bogus accounts, as long as there are only a few certificates from the "good" web to the bogus accounts. Those certificates represent a bottleneck in the network flow.

Mapping into graph

The mapping of certificates into a graph is dependent on a parameter: the certification level l . Each account on Advogato corresponds to a node in the graph. An edge exists from node s to node t when account s has certified account t at level l or higher. In addition, there is a distinguished "seed" node, with predefined edges to accounts.

Security proof

The nodes are split into three categories: good, confused, and bad. The bad nodes are under the attacker's control. The confused nodes themselves represent valid accounts, but may contain certificates to the bad nodes. The good nodes are both valid accounts and have certificates only for other good nodes and confused nodes. This partition is shown graphically below:

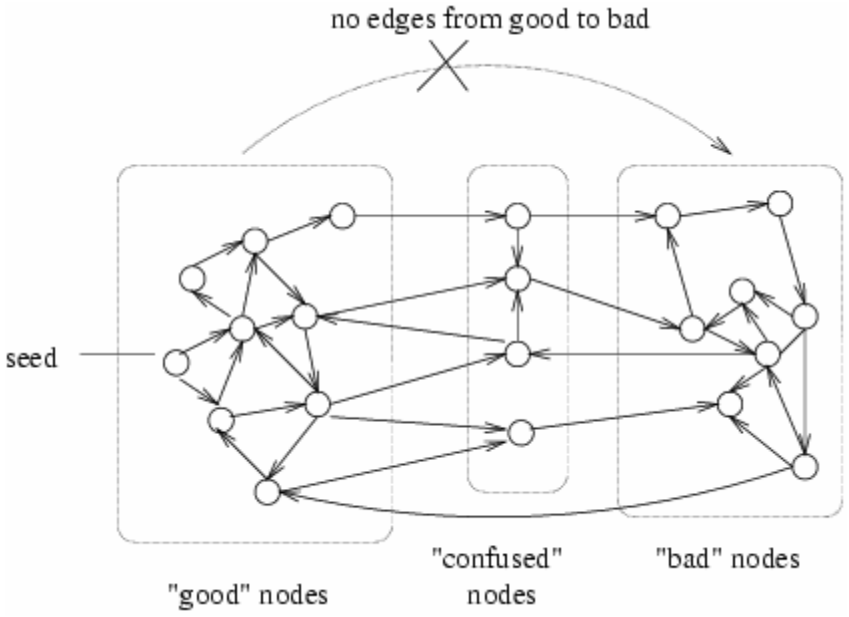

Conclusion

The trust metric used in Advogato has a property not known in any previous trust metric: resistance to catastrophic failure in the face of a sufficiently massive attack. Instead, the number of bad nodes accepted scales linearly, and with a fairly small constant, with the number of certificates from valid accounts to bogus ones. It is also easy to compute efficiently and fairly simple to understand. As such, it should find applications in security infrastructures, as well as defining online communities, reliably excluding spammers, trolls, and other common annoyances.

REFERENCES

Ager, David L. and Mikolaj J. Piskorski. 2001. "Age, Audiences and Autonomy in the Dynamics of Organizational Prestige." in *Working Paper*. Harvard Business School.

Asch, Morton Jay. 1951. "Effects of Group Pressure Upon the Modification and Distortion of Judgments." in *Groups, Leadership, and Men*, edited by H. Guetzkow.

Axelrod, Robert M. and Michael D. Cohen. 1999. *Harnessing complexity : organizational implications of a scientific frontier*. New York: Free Press.

Baker, W. E. and R. R. Faulkner. 1991. "Role as Resource in the Hollywood Film Industry." *American Journal of Sociology* 97:279-309.

Bearman, Peter. 1997. "Generalized Exchange." *American Journal of Sociology* 102:1383-1415.

Berger, Joseph, David G. Wagner, and Morris Zelditch. 1992. "A Working Strategy for Constructing Theories: State Organizing Processes." Pp. 107-123 in *Metatheorizing*, edited by G. Ritzer. Newbury Park, CA: Sage.

Berger, Joseph and Morris Zelditch. 1998. *Status, power, and legitimacy : strategies and theories*. New Brunswick, NJ: Transaction.

Berger, Peter L. and Thomas Luckmann. 1966. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. New York: Doubleday & Company, Inc.

Blau, Peter Michael. 1964. *Exchange and power in social life*. New York: Wiley.

Blossfeld, Hans-Peter and Götz Rohwer. 1995. *Techniques of event history modeling : new approaches to causal analysis*. Mahwah, NJ: L. Erlbaum.

Chalmers, Rachel. 2000. "Even Better than Slashdot?" in *Salon.com*.

Davis, G. F. and H. R. Greve. 1997. "Corporate elite networks and governance changes in the 1980s." *American Journal of Sociology* 103:1-37.

Dimaggio, P. J. and W. W. Powell. 1983. "The Iron Cage Revisited : Institutional Isomorphism and Collective Rationality in Organizational Fields." *American Sociological Review* 48:147-160.

Festinger, Leon. 1957. *A theory of cognitive dissonance*. Evanston, IL: Row Peterson.

Fombrun, Charles J. 2001. "Corporate Reputations as Economic Assets." in *The Blackwell Handbook of Strategic Management*, edited by M. A. Hitt, R. E. Freeman, and J. S. Harrison. Oxford: Blackwell Publishers Ltd.

Frank, Robert H. 1985. *Choosing the right pond : human behavior and the quest for status*. New York ; Oxford [Oxfordshire]: Oxford University Press.

Ghosh, Rishab Aiyer. 1998. "FM Interview with Linus Torvalds: What Motivates Free Software Developers?" *First Monday* 3.

Goode, William J. 1978. *The Celebration of Heroes: Prestige as a Control System*. Berkeley: University of California Press.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360-1380.

—. 1974. *Getting a job; a study of contacts and careers*. Cambridge, Mass., Harvard University Press,.

Greve, H. R. 1996. "Patterns of competition: The diffusion of a market position in radio broadcasting." *Administrative Science Quarterly* 41:29-60.

Haunschild, P. R. 1994. "How Much Is That Company Worth : Interorganizational Relationships ; Uncertainty ; and Acquisition Premiums." *Administrative Science Quarterly* 39:391-411.

Heider, Fritz. 1958. *The psychology of interpersonal relations*. New York,: Wiley.

Homans, George Caspar. 1950. *The human group*. New York, Harcourt, Brace,.

Kollock, P. 1994. "The Emergence of Exchange Structures : An Experimental-Study of Uncertainty ; Commitment ; and Trust." *American Journal of Sociology* 100:313-345.

Leifer, Eric M. and H. C. White. 1987. "A Structural Approach to Markets." Pp. 85-108 in *Intercorporate Relations: The Structural Analysis of Business*, edited by M. S. Mizruchi and M. Schwartz. Cambridge: Cambridge University Press.

Lerner, Josh and Jean Tirole. 2002. "The Simple Economics of Open Source." *Journal of Industrial Economics* 52.

Merton, Robert King. 1968. "The Matthew Effect in Science." *Science* 159:56-63.

Meyer, J. W. and B. Rowan. 1977. "Institutionalized Organizations : Formal-Structure as Myth and Ceremony." *American Journal of Sociology* 83:340-363.

Milgram, Stanley. 1967. "The Small-World Problem." *Psychology Today* 1:60-67.

Milgrom, Paul R. and John Roberts. 1992. *Economics, organization, and management*. Englewood Cliffs, NJ: Prentice-Hall.

Mockus, Audris, Roy Fielding, and James Herbsled. 2000. "A Case Study of Open Source Software Development: The Apache Server." in *Proceedings of the International Conference on Software Engineering*. Limerick, Ireland.

Newcomb, Theodore Mead. 1943. *Personality and social change; attitude formation in a student community*. New York: Dryden Press.

Philips, Damon J. and Ezra W. Zuckerman. 2001. "Middle-Status Conformity: Theoretical Restatement and Empirical Demonstration in Two Markets." *American Journal of Sociology* 107:379-429.

Podolny, J. M. 1993. "A Status-Based Model of Market Competition." *American Journal of Sociology* 98:829-872.

—. 2001. "Networks as the Pipes and Prisms of the Market." *American Journal of Sociology* 107:33-60.

Podolny, J. M., T. E. Stuart, and M. T. Hannan. 1996. "Networks ; knowledge ; and niches: Competition in the worldwide semiconductor industry ; 1984-1991." *American Journal of Sociology* 102:659-689.

Rangan, S. 2000. "The problem of search and deliberation in economic action: When social networks really matter." *Academy of Management Review* 25:813-828.

Raymond, Eric S. 1999. *The Cathedral and the Bazaar*. Sebastopol: O'Reilly & Associates.

Ridgeway, C. L. and K. G. Erickson. 2000. "Creating and spreading status beliefs." *American Journal of Sociology* 106:579-615.

Ross, Lee and Richard E. Nisbett. 1991. *The person and the situation : perspectives of social psychology*. Philadelphia, Pa.: Temple University Press.

Sherif, Muzafer. 1966. *In common predicament; social psychology of intergroup conflict and cooperation*. Boston,: Houghton Mifflin.

Spence, A. Michael. 1974. *Market Signaling: Informational Transfer in Hiring and Related Processes*. Cambridge, MA: Harvard University Press.

Stallman, R. 1999. "The GNU Operating System and the Free Software Movement." in *Open Sources*, edited by C. Dibona, S. Ockman, and M. Stone. Sebastopol, CA: O'Reilly.

Stuart, T. E., H. Hoang, and R. C. Hybels. 1999. "Interorganizational endorsements and the performance of entrepreneurial ventures." *Administrative Science Quarterly* 44:315-349.

Takahashi, Nobuyuki. 2000. "The Emergence of Generalized Exchange." *American Journal of Sociology* 105:1105-1134.

Thye, S. R. 2000. "A status value theory of power in exchange relations." *American Sociological Review* 65:407-432.

Torrance, E. Paul. 1955. "Some Consequences of Power Differences on Decision Making in Permanent and Temporary Three-Man Groups." Pp. 482-492 in *Small Groups*, edited by A. P. Hare, E. F. Borgatta, and R. F. Bales. New York: Knopf.

Tuma, Nancy B. 1980. *Invoking RATE*. Mannheim: ZUMA.

Watts, D. J. 1999. "Networks ; dynamics ; and the small-world phenomenon." *American Journal of Sociology* 105:493-527.

Weber, Max. 1968. *Economy and society; an outline of interpretive sociology*. New York,: Bedminster Press.

White, H. C. 1981. "Where Do Markets Come From." *American Journal of Sociology* 87:517-547.

Zuckerman, E. W. 1999. "The categorical imperative: Securities analysts and the illegitimacy discount." *American Journal of Sociology* 104:1398-1438.