

Heterogeneity, Brokerage, and Innovative Performance: Endogenous Formation of Collaborative Inventor Networks

Jeongsik “Jay” Lee

College of Management, Georgia Institute of Technology, Atlanta, Georgia 30308,
jay.lee@mgt.gatech.edu

In this study, I examine how past performance influences the relative positions of actors in a network and how the actor-level heterogeneity in quality mediates the often-demonstrated association between brokerage and performance. On the collaboration network of U.S. biotech inventors during 1976–1995, I find that inventors with superior track records are more apt to form collaboration ties that enhance brokerage, thereby occupying positions that allow them to broker across network boundaries. Controlling for past performance significantly weakens the positive relationship between brokering position and innovative performance. Furthermore, when inventor-level heterogeneity is controlled for through inventor fixed effects, the position-performance correlation disappears. These findings suggest that, at least for collaborative inventors, actor-level heterogeneity such as performance history largely drives the asymmetry in brokerage, explaining most of the position–performance association.

Key words: social networks; innovation; brokerage; heterogeneity; performance; endogeneity; mediation effect

History: Published online in *Articles in Advance* December 4, 2009.

1. Introduction

Social networks research provides ample evidence that actors' relative positions in a network correlate with their economic performance. Actors occupying valuable network positions—though researchers agree less on which positions qualify as valuable—are shown to outperform those who do not. To cite only a few, managers linked to others who are disconnected from each other—i.e., those that occupy brokering positions—are promoted faster, generate better ideas, and receive more-favorable evaluations (Burt 1997, 2004; Podolny and Baron 1997); engineers situated in dense networks show higher involvement in innovation (Obstfeld 2005); and research and development R&D teams of scientists closely connected to the team members while maintaining diverse outside contacts are more productive (Reagans and Zuckerman 2001).

Although this literature has elucidated how differences in network positions relate to performance differentials among actors, it has given little attention to how the heterogeneity in actor attributes—such as performance history—generates this positional asymmetry. For instance, scholars have demonstrated that actors in brokering positions that provide the sole links between others exhibit superior performance. This occurs primarily through preferential access to information sources and through control of information flow (Burt 1992, 1997, 2004; Hargadon and Sutton 1997; Podolny and Baron 1997; Zaheer and Bell 2005; Shipilov 2006). However, it is likely that better-performing actors came to occupy brokering positions through their ability to secure the

opportunities leading to those positions. This happens if actors with superior quality—which also generates performance excellence—can identify and exploit opportunities to form nonredundant ties more efficiently, or if their perceived quality based on past performance attracts others to confer such connections on them, or both. Research provides circumstantial evidence of these phenomena. Differentials in underlying quality give rise to different propensities for forming network relationships. At the individual level, high-quality scientists not only publish more but also are more apt to form collaborative relationships (Zuckerman 1967, Zucker and Darby 2006). At the organizational level, it is the better-performing firms that enter into strategic alliances (Chan et al. 1997). Recognition, reputation, or awards accrue to actors exhibiting exceptional performance, but often disproportionately to their underlying quality (Merton 1968, 1988). Actors form beliefs about the quality of others based on revealed performance. Thus, excellence in past performance—even if it is owed partly to luck—can lead to advantages in forming network ties that in turn generate higher returns. For network brokers, these will be the opportunities to establish new, nonredundant ties (Burt 1992). The advantages arising from performance history tend to get amplified due to this signaling effect (Gould 2002).

Therefore, in the context of brokerage, actors who exhibit superior performance come to occupy seemingly valuable network positions by being able to form network ties that lead them to such positions, rather than the reverse (favorable positions enabling

their occupants to outperform those situated less favorably). This implies that brokering positions arise endogenously. Actor-level heterogeneity that generates performance differential also influences network tie formation, mediating the position–performance association. Then, the often-demonstrated effect of brokering position on performance could have been confounded by this performance–position feedback underlying it. Hence, I expect that (i) actors' past performance will be positively associated with the probability that they form collaboration ties that lead to brokering positions, and (ii) the actor-level heterogeneity in performance positively mediates the brokering position–performance relationship; that is, if the actor-level heterogeneity is accounted for, the positive association between brokering position and performance will significantly weaken.

I test these hypotheses on the collaboration network of inventors in the U.S. biotechnology industry during the period of 1976–1995. I find that (i) high-performing inventors are more apt to form collaboration ties that increase brokerage, after controlling for other factors such as dyad attributes (e.g., past coauthorship) and tie characteristics (e.g., size of ego network); (ii) without controls for inventor-level heterogeneity in performance, brokering position correlates positively with innovative performance, consistent with findings in prior studies (e.g., Burt 2004); (iii) controlling for past performance significantly reduces the strength of the brokering position–patent performance association; and (iv) when inventor-level heterogeneity is controlled for through inventor fixed effects, the position–performance correlation disappears. These results are robust across multiple measures of brokerage and performance and across different specification methods. At least in the collaborative network of biotech inventors, actor-level heterogeneity such as performance history largely determines the positional asymmetry, and the returns attributable to brokerage are much lower than they appear in cross-section studies.

2. Actor Heterogeneity, Network Position, and Performance

2.1. Network Position as Correlate of Performance

Scholars of social networks have consistently demonstrated a significant association between network position and performance. Actors' positions in an exchange network—i.e., how they are linked to others—have been shown to correlate with economic performance. However, researchers have been equivocal as to which type of network position leads to performance excellence. One line of research indicates that a cohesive network structure, in which actors are densely linked to each other, engenders mutual trust, promotes individual attachment to the group, and reduces barriers to resource mobilization, thereby enhancing performance (Ahuja 2000,

Fernandez et al. 2000, Obstfeld 2005). Other research emphasizes the benefit of brokerage. Actors brokering between otherwise disconnected actors are in an advantageous position for identifying arbitrage opportunities, have higher chances of creating new knowledge or products, and are better able to capitalize on their existing capabilities (Burt 1997, 2004; Hargadon and Sutton 1997; Zaheer and Bell 2005; Shipilov 2006). The contingent benefit of both types of network positions has also been suggested (Reagans and McEvily 2003, Soda et al. 2004, Fleming et al. 2007). The differences in perspective notwithstanding, the consensus has been that network positions correlate significantly with actor performance.

In contrast with the abundant evidence of position–performance correlation, we have very limited understanding of how actors come to occupy certain positions. Researchers have proposed theoretical models that characterize the evolutionary process of network formation.¹ However, focusing on the efficiency of equilibrium network structure, these models provide limited insight into how the relative positions of actors are determined. On the empirical side, researchers have examined the correlates of network activities such as tie formation. Among them are prior network experience (Powell et al. 1996), network position or status (Gulati and Gargiulo 1999, Sorenson and Stuart 2001, Powell et al. 2005), other actors' performance history (Baum et al. 2005), and resource dependence on others (Gulati and Gargiulo 1999). However, when explaining the association between network activity and performance, researchers implicitly assume that network positions are given. As such, the research focus has predominantly been to examine how the existing distribution of positions helps explain the performance asymmetry among actors. Without understanding the mechanism that generates the distribution of network positions, however, one cannot determine whether performance advantage stems from the position, or from other factors that curb competition. Thus, uncovering the factors behind the positional asymmetry becomes imperative for better understanding the role of network positions in shaping economic outcomes. In this study, I examine the generative process of positional asymmetry and propose performance heterogeneity as an important antecedent of such asymmetry. I also explore the implication of the feedback from past performance to network formation on the position–performance association. In discussing how performance differentials drive actors to occupy different network positions, I focus on brokerage as a characterization of actor position.

2.2. Brokerage as Network Position

Brokerage is widely used in the social networks literature for characterizing actor position in a network (Burt 1997, 2004; Hargadon and Sutton 1997; Ahuja 2000;

Zaheer and Bell 2005). Actors are in a brokering position when they provide the only connection between two actors. Thus, any flow of information, knowledge, or products from one actor to another necessarily passes through the brokering actor. In Burt's (1992) terminology, structural holes exist between these brokered actors. Hence, actors in a brokering position—thereby spanning the structural holes around them—can enjoy the benefit of preferential access to information and, if necessary, skew the flow of information to their advantage.

Considering the information benefit of brokerage, which emphasizes the novelty of ideas and information flowing from nonredundant ties (e.g., Burt 2004), it seems natural to examine innovation as a correlate of brokering position. Innovation is a knowledge-intensive activity, which builds upon identifying, acquiring, assimilating, and combining knowledge (Cohen and Levinthal 1990, Rosenkopf and Nerkar 2001, Fleming 2001). Hence, the extent to which actors secure access to diverse knowledge sources and exploit the informational advantage determines the efficacy of their network positions and their ability to achieve innovative excellence. Many studies on innovation have focused on brokering position in knowledge networks (Reagans and McEvily 2003, Burt 2004, Nerkar and Paruchuri 2005). Brokerage is also ideal for characterizing network position particularly given the empirical context of this study—biotechnology, a highly knowledge-intensive and innovation-driven scientific field.

The value of a brokering position stems from the nonredundancy of network ties. Actors hardly benefit from adding links to an actor who is already connected through other ties. No new information becomes available through this additional tie. Moreover, the focal actor cannot hope to enjoy an advantage by controlling the flow of information, because alternative routes for such flow are available to the other party. New ideas and opportunities only come through nonredundant ties that bridge over structural holes (Burt 1992). Actors who maintain less redundant network ties perform more valuable brokering roles in the network and can appropriate more from their positions. In brokerage, therefore, establishing new, nonredundant ties is critical for achieving performance excellence. Below, I discuss how the brokering position can be determined by actor-level attributes, particularly past performance.

2.3. Performance Heterogeneity and Development of Brokering Position

Networks are complex constructs, with numerous factors contributing to their generation. Thus, actor performance may also function as an antecedent to networking activities, rather than only being an output. In fact, research finds that actors consider their potential partner's performance record when making tie formation decisions (Baum et al. 2005). There are two possible conduits by

which performance relates to the formation of ties leading to a brokering position: performance as a representation of underlying quality and performance as a signal.

2.3.1. Performance Representing Underlying Quality. Actors are heterogeneous in their quality.² In particular, actors differ in their ability to perform given tasks. Unless outcomes are shaped purely by random chance, the underlying quality of actors determines their performance on the tasks. Put differently, higher-quality actors are likely to exhibit superior performance. Performing a task may involve collaboration with other actors. In general, the costs and benefits associated with collaboration differ across types of ties; certain types of collaboration ties are more costly to form but generate higher returns. Then, high-performing actors will likely choose to form ties with greater returns because the opportunity cost they incur for such collaboration is higher than that of low-performing counterparts.³

In the context of brokerage, actors exhibiting high performance may be those occupying network positions that broker across actors (Burt 1992). Collaborating with a new and structurally distant other is costly and involves uncertainty. It takes resources such as time and energy to search for candidates, initiate the contact, and maintain the relationship. The search may extend beyond the boundary of primary activities, imposing extra cost to the initiator. The newly formed tie may turn out to be unsatisfactory, in which case one has to incur additional cost to repeat the process. However, according to the brokerage argument (e.g., Burt 1992), the expected returns from initiating new collaboration ties are also higher because nonredundant ties bring the greatest network benefits. In contrast, repeated collaboration is less costly, if not costless. Information asymmetry is less of a concern. Familiarity may reduce possible conflicts. Certain routines may have been developed, improving the collective efficiency. However, the expected returns are also lower because new ideas or opportunities for improvement seldom arise from such a relationship.⁴ In sum, high-quality actors are apt to achieve high performance owing to their superior quality. At the same time, actors of higher quality will form more new, nonredundant ties relative to their lower-quality counterparts, given the opportunity cost of collaboration. Therefore, actors with better performance records have a higher likelihood of occupying brokering positions—characterized by many nonredundant ties around them—owing to their higher quality.

2.3.2. Performance as Signal. Performance, on the other hand, functions as a signal to potential ties. Forming a collaborative relationship with others involves uncertainty. In general, this uncertainty concerns two aspects: whether the other actor has the ability to perform the task and how well the actor can do so. The uncertainty associated with a tie formation decision is

lower if more information is available on the other actor. Everything else equal, actors prefer those who have some performance record, and this preference strengthens if the quality of work is difficult to verify. A performance record indicates an actor's ability, reducing the uncertainty. A good performance record signals the superiority in underlying quality, and hence increases the likelihood that the actor will be chosen as a tie.

Good performance also expands the opportunities for being considered as a potential tie. Actors with proven performance come to occupy positions in which they become more visible to potential exchange partners. Recognition, reputation, or awards are given to actors that perform well. To the extent that these connections in turn contribute to future performance, this dynamic would produce a virtuous cycle between position and performance, amplifying the quality differences across actors (Gould 2002). This cumulative advantage accrues to actors who initially exhibit superior performance, conferring on them disproportionately more opportunities to form new ties.⁵

This signaling effect is likely to be the strongest for new ties. For existing ties, revealed performance is hardly news. For potential exchange partners, however, performance records can be a critical, and often the only, piece of information available for tie formation decisions. Thus, opportunities for establishing new ties accrue more to proven actors. Furthermore, actors receiving more offers are able to select among alternatives. Everything else equal, they will choose ties of the greatest returns. Hence, actors with superior track records will form more new collaboration ties.

To summarize, actor performance represents the underlying quality that determines the extent to which an actor forms new collaboration ties. It also signals the actor's potential value as an exchange partner, thereby expanding the opportunity set of tie formation. In either case, past performance will be a significant correlate of tie formation pattern, with a high-performance actor exhibiting a higher probability of forming collaboration ties that lead to a brokering position. The discussions so far lead us to the following hypothesis:

HYPOTHESIS 1. An actor's past performance will be positively associated with the likelihood of forming a collaboration tie that increases brokerage.

2.4. Actor-Level Heterogeneity Mediates Brokering Position–Performance Association

I have argued that actors are heterogeneous in their ability to perform certain tasks, with high-quality actors exhibiting higher level of performance in the relevant task domain. Indeed, scholars have found that the distribution of revealed performance is far from being uniform (e.g., Lotka 1926), and high-quality actors are much more productive than their average peers (e.g.,

Zuckerman 1967). Moreover, such performance differential tends to persist (e.g., Gompers et al. 2006). The evidence collectively suggests that actor-level heterogeneity in quality strongly and positively correlates with performance. I have also argued that high-performing actors are more likely to occupy brokering positions. Put differently, actor quality positively correlates with the probability of an actor's arriving at a brokering position.

The existence of positive links from actor-level heterogeneity to brokering position and performance suggests that actor quality essentially mediates the often-demonstrated positive association between brokering position and performance in the social networks literature (e.g., Burt 1997, 2004; Hargadon and Sutton 1997; Zaheer and Bell 2005; Shipilov 2006). Although scholars have repeatedly shown that actors occupying brokering positions enjoy performance advantages, the causal mechanisms that underlie such association have often been taken for granted rather than explicitly modeled (cf. Reagans and McEvily 2003). Hence, it appears less clear whether there really is an association between brokering position and performance, particularly given the argued mediation by actor-level heterogeneity.

Mediation occurs when certain construct intervenes between two other constructs, reinforcing or even generating a seemingly causal relationship between the two (Baron and Kenny 1986). In the presence of mediation, if the researcher appropriately controls for this indirect effect through the mediator, the association between two constructs is either significantly reduced in magnitude—when the mediation is moderate—or completely disappears—when there is a perfect mediation. Given the mediating role of actor-level performance heterogeneity, brokering positions may appear valuable because actors with superior performance come to occupy those positions, not necessarily because the positions enable them to outperform others. Empirically, the mediation of performance heterogeneity implies that, without controls for actor-level heterogeneity in past performance, brokering position may correlate positively with future performance; however, when such heterogeneity is accounted for, the association between brokering position and future performance will significantly weaken. Failures to account for the effect from actor-level heterogeneity can thus lead to an overestimation of the association between brokering position and performance. Hence, I hypothesize the following:

HYPOTHESIS 2. Actor-level heterogeneity in past performance positively mediates the relationship between brokering position and future performance.

3. Empirical Models

3.1. Overview

My empirical corroboration consists of two parts. The first part examines how past performance relates to the

pattern of tie formation. For this, I estimate the coefficient on the past performance measure in relating to the probability that, conditional upon a collaboration tie being formed, the given tie represents a collaborative relationship that increases brokerage. The second part investigates the degree to which actor-level heterogeneity in quality mediates the brokering position–performance association. In this analysis, I estimate and compare the coefficients on the brokerage measure across models with varying controls for actor-level heterogeneity.

The empirical context of this study is the patent collaboration among inventors in the U.S. biotechnology industry from 1976 to 1995. Although not all inventions are patentable and the propensity to patent differs across patenting agencies, patents are widely used to measure knowledge and technological innovation (e.g., Griliches 1990, Lanjouw and Schankerman 2004). The biotechnology industry is one of the most innovative and knowledge-intensive fields. Collaboration between and within organizations is active and commonplace (Zucker et al. 1998, Nerkar and Paruchuri 2005, Powell et al. 2005). Furthermore, patents are a highly effective means of appropriation and are used consistently across firms (Cohen et al. 2000). Hence, patent collaboration among biotech researchers provides a rich setting for observing the evolution of networks and testing the relationship between network position and innovation performance.

3.2. Data

The primary source of the data is the National Bureau of Economic Research (NBER) patent data set (Hall et al. 2001), which was later extended by Hall.⁶ The data set includes all U.S. utility patents granted during 1963–2002. I chose the period 1976–1995 for the sample construction. The biotech industry did not exist until the 1973 discovery of recombinant DNA (Zucker et al. 1998). Thus, any serious collaboration and the resulting innovation occurred after this event. Also, it may take a few years to obtain patentable outcomes from any project. Hence, I began the sample with patents applied for from 1976.⁷ A measure of innovation performance in this study requires sufficient time lags between grant date and citation to minimize the truncation bias. Research suggests that most citations are made within five years from patent grant (e.g., Lanjouw and Schankerman 2004). Furthermore, it takes about two years from application to grant.⁸ Hence, I ended the sample with patents applied for by 1995.

I defined the network to include every inventor who has made at least one successful U.S. patent application in biotechnology during the sample period. To identify the biotech patents, I followed the definition by the United States Patent and Trademark Office (2002; USPTO). The USPTO defines biotechnology based on the patent classes. If a patent's primary class belongs to

biotechnology, I classified it as a biotech patent. This criterion identified 26,330 U.S. biotech patents. To identify the distinct inventors associated with these patents, I mainly relied on the paper by Trajtenberg et al. (2006), who developed a “who is who” in the NBER inventor data, and made further adjustments.⁹ The appendix details the inventor identification procedure. This process identified 28,267 unique inventors. These inventors also have 31,275 nonbiotech patents.¹⁰ Hence, the final data set contains 57,605 U.S. patents by 28,267 biotech inventors.

3.3. Analysis 1: Collaboration Tie Formation

The first analysis concerns the correlates of tie formation. Conditional on collaboration, three types of tie formation are possible: collaborating with an existing tie (I label this a “repeating tie”), collaborating with an inventor new to the focal inventor but linked to existing collaborators (a “redundant tie”), and collaborating with an inventor new to everyone in the ego network (a “nonredundant tie”). During 1981–1995, inventors formed 200,515 collaboration ties.¹¹ I used a five-year window to define the network relationship (e.g., Sorenson and Stuart 2001). Among the ties, 68,564 (34.2%) were repeating ties, 4,924 (2.5%) were redundant ties, and 127,027 (63.4%) were nonredundant ties.

Various factors could be associated with the probability of tie formation between inventors. Not only do the focal inventor's attributes affect the tie formation, but tie characteristics and dyad attributes can also influence this likelihood. The case-control method is an appropriate empirical design that controls for different risks of collaboration due to these heterogeneous attributes (Sorenson and Stuart 2001, Powell et al. 2005). I used a stratified sampling scheme to ensure that the resulting sample mimics the network structure of inventors while containing a sufficient number of controls for types that, when randomly sampled, would be underrepresented. Specifically, for each realized tie, I randomly selected from each type one control that could have been realized but was not. Hence, a realized tie at time t was, in principle, matched with three controls (one *repeating tie*, one *redundant tie*, and one *nonredundant tie*). Ties lacking alternatives of a certain type were included without a matched control of that type. This process yielded a sample of 569,872 observations with 200,515 cases and 369,357 controls.¹²

The three types of ties together determine brokering position. Not all types, however, exert the same influence on inventor position. If the value of brokerage stems from the nonredundancy of ties (Burt 1992), one's brokering position is best enhanced by the formation of a nonredundant tie. Repeating ties do not affect the position,¹³ and collaboration with a redundant tie might decrease brokerage.¹⁴ Hence, I focus on nonredundant

tie formation. Specifically, I test whether past performance positively correlates with the probability that, contingent upon a collaboration tie occurring, a given tie is nonredundant, after controlling for attributes of the tie and the dyad. A significant coefficient on the past performance measure would imply that past performance influences an inventor's brokering position in the collaborative network.

3.3.1. Measures. Three sets of explanatory variables were included in the model: focal inventor attributes (patent performance, network position, and number of collaborators), tie attributes (new inventor, number of collaborators, and patent performance), and dyad attributes (prior ties, coauthorship, organizational affiliation, and geographical location).

Patent Performance. *Simple patent counts* and *citation-weighted patent counts* were used as performance measures. The first measure counts the number of patents an inventor successfully applied for each year. A patent with multiple inventors was counted as $1/N$ toward each inventor, where N is the number of inventors collaborated on the patent.¹⁵ Insofar as a patent signifies novelty (Griliches 1990), this variable measures the amount of novel knowledge created. The second measure counts all citations made to each patent, divided by the number of associated inventors. Hence, an inventor who received no citations was assigned zero for the variable. Researchers have used the number of citations to measure the importance or quality of the patent (Trajtenberg 1990, Lanjouw and Schankerman 2004, Hall et al. 2005). To the extent that patent citations indicate innovation quality, this variable measures the quality-adjusted knowledge created.

For measures of past performance, I constructed *patent stock* by cumulating the number of patents or citations up to period $t - 1$, where t is the year of tie formation.¹⁶

Brokering Position. I used network constraint (Burt 1992) to measure *brokerage*.¹⁷ The formula for this measure is

$$C_i = \sum_j \left(P_{ij} + \sum_{q \neq i \neq j} P_{iq} P_{qj} \right)^2,$$

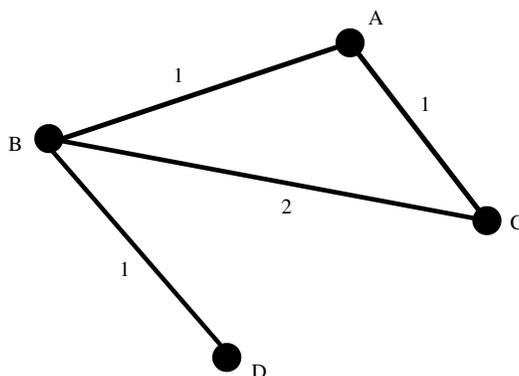
where C_i is the network constraint of inventor i , and p_{ij} is the proportion of time and energy that inventor i invested in inventor j (p_{iq} and p_{qj} are defined analogously).¹⁸ Hence, C_i is the proportion of inventor i 's relationship that are invested in connection with inventor j . Lower values on this measure imply that the inventors occupy less constrained positions, thereby brokering more extensively in the network. Figure 1 illustrates how the network constraint was computed from the inventor collaboration network data. In this subnetwork, there are four inventors, A through D. During a five-year period, A has collaborated on one patent with

B and C. During the same period, B and C have collaborated on two patents, one of which with A. B also has another patent in collaboration with D. Applying the formula above, network constraints for A and B are computed as 1.257 and 0.627, respectively. Clearly, B is much less constrained than A, and hence brokers more widely across inventors. For an easier interpretation, I transformed the network constraint by subtracting it from 2 so that higher values indicate higher brokerage.¹⁹ For inventors without collaborators, I assigned zero to *brokerage*. To incorporate the eroding network relationships over time (Burt 2000), I used a five-year moving window of collaboration network, beginning from 1976. In the model specification, I lagged the measure by one year to avoid simultaneity. Current brokering position was included to control for possible effects from occupying certain positions.

Coauthorship. In biotech, many scientific discoveries published in academic journals also get patented. Murray and Stern (2007) report that almost half of the articles in their sample resulted in a U.S. patent grant. Thus, coauthorships in scientific articles are likely to entail collaboration in patenting. Moreover, to the extent that these two knowledge production networks overlap, article coauthorships will also affect the pattern of tie formation in collaborative patenting. Hence, I included a *coauthorship* dummy in the model to control for the effects of this prior relationship on patent collaboration and the type of ties that are formed.²⁰ This dummy indicates if an inventor has coauthored a scientific article with the other inventor prior to year t . For identifying prior coauthorship, I searched the Institute for Scientific Information (ISI) Web of Science database using the last name and the initials of first name and middle name of inventors.²¹ This process identified a total of 966,797 articles by 22,784 (80.6%) inventors published during 1974–1995.²² Of these inventors, 11,190 (39.6%) have coauthored with at least one other inventor.

Other Controls. *Number of collaborators* measures the network size of an inventor and was included to control for the size effect on the choice of collaboration ties. If tie formation follows a preferential attachment, inventors with larger networks may be more likely to form a nonredundant tie. The converse is likely if there are negative returns to scale. Moreover, this variable controls for differences in the risk of repeating ties. I included the variable for both focal inventor and tie. The measure counts the number of distinct collaborators in each of the five-year ego networks. *Patent stock* of the tie controls for the possible influence of a partner's experience on tie formation. The *new inventor* dummy indicates if the inventor patented for the first time in that year, and was included to control for its asymmetric effects on the probability of forming certain types of ties. *Repeating tie* indicates if the focal inventor had a prior copatenting relationship with the tie. Similarly,

Figure 1 Illustration of Inventor Collaboration Network (Subcomponent) and Computation of Network Constraint Measures



	A		B	
Contact-specific constraints	B	0.694 [= (1/2 + 1/2 * 2/3) ²]	A	0.174 [= (1/4 + (2/4 * 1/3 + 1/4 * 0/1)) ²]
	C	0.563 [= (1/2 + 1/2 * 2/4) ²]	C	0.391 [= (2/4 + (1/4 * 1/2 + 1/4 * 0/1)) ²]
			D	0.063 [= (1/4 + (1/4 * 0/2 + 2/4 * 0/3)) ²]
Aggregate constraint	Total	1.257	Total	0.627

redundant tie indicates if the tie has copatented with one of the focal inventor’s collaborators. Everything else equal, inventors are more likely to collaborate with those whom they have prior relationships with. The *coaffiliation* dummy controls for differences in the risk of collaborating due to organizational affiliation.²³ Similarly, the *colocation* dummy controls for differences in the collaboration probability due to geographic distance.²⁴ Finally, *year* dummies were included to control for time fixed effects. Table 1 provides descriptive statistics of the variables used for specification. The final sample included 183,577 cases and 346,883 controls by 24,816 inventors.²⁵

3.3.2. *Estimation Method.* The dependent variable is a binary indicator of forming a collaboration tie. With a random sample, a standard logit model would consistently estimate the coefficients. Sampled on the dependent variable, however, each type of tie has different probability of entering the sample. The maximum likelihood estimator devised for random samples is generally inconsistent when applied to a choice-based sample such as this. Manski and Lerman (1977) proposed a maximum likelihood estimator that weights each observation’s contribution to the log-likelihood, thereby producing consistent estimates for a choice-based sample.²⁶ This weighted exogenous sample maximum likelihood (WESML) estimator is also asymptotically efficient (Cosslett 1981). Management scholars have applied this estimation method to analyzing samples constructed following this procedure (e.g., Singh 2005). For the weights, I assigned to the controls an inverse of the probability that each type was sampled. To each case, I assigned “1” because all realized ties were

included. In essence, this treatment gives greater weights to the observations that are undersampled such that the weighted sample simulates the underlying distribution of the population.

It should be noted that an inventor’s forming a collaboration tie at time *t* is conditioned on the inventor’s producing at least one collaborative patent at *t*.²⁷ Not every inventor has the same ability or propensity to file for a collaborative patent at a given period. Failures to account for the differences in the patenting hazard may introduce bias to the estimates. Thus, I first ran a probit model and included the *nonselection hazard*—i.e., the inverse Mill’s ratio—as a regressor in the tie formation equation (Heckman 1979).²⁸ The dependent variable in the selection equation takes 1 if the inventor had a collaborative patent at *t* and 0 otherwise.²⁹ For the explanatory variables, I used *brokerage*, the *number of collaborators*, *patent stock*, the *size of inventor pool*, and the *proportion of coauthored scientific articles*. The last variable—computed as the number of collaborated articles divided by the total number of articles authored by the inventor—captures the proclivity to collaborate on article authorship. These explanatory variables were all lagged by one year to avoid simultaneity. I also included *year* dummies and a dummy for the first-time inventors who thus have no *brokerage* measure—assigned zero—for the period. Because the *nonselection hazard* term included in the tie formation model is estimated with error, the standard errors of the WESML coefficients may be biased, causing an overrejection of the null hypothesis of noneffect (cf. Hamilton and Nickerson 2003). To correct for this problem, I bootstrapped the standard errors instead of simply reporting the errors

Table 1 Descriptive Statistics (Analysis of Collaboration Tie Formation)

Variables	All (N = 530,460)				Cases (N = 183,577)		Controls (N = 346,883)	
	Mean	Std. dev.	Min	Max	Mean	Std. dev.	Mean	Std. dev.
Ln(Patent stock $t - 1$) (focal)	0.99	0.77	0.05	5.34	0.85	0.73	1.06	0.78
Ln(CW patent stock $t - 1$)	2.34	1.49	0	8.65	2.09	1.48	2.47	1.48
Brokerage	1.08	0.66	0	1.96	0.91	0.70	1.18	0.61
Ln(No. collaborators) (focal)	1.53	1.05	0	4.29	1.26	1.08	1.68	1.01
Repeating tie dummy	0.27	0.44	0	1	0.37	0.48	0.21	0.41
Redundant tie dummy	0.19	0.39	0	1	0.03	0.16	0.27	0.44
Nonredundant tie dummy	0.54	0.50	0	1	0.60	0.49	0.51	0.50
Coauthorship dummy	0.14	0.34	0	1	0.25	0.43	0.08	0.27
Coaffiliation dummy	0.36	0.48	0	1	0.88	0.32	0.08	0.28
Colocation dummy	0.39	0.49	0	1	0.79	0.41	0.17	0.38
New inventor dummy (tie)	0.30	0.46	0	1	0.42	0.49	0.24	0.42
Ln(No. collaborators) (tie)	1.20	0.95	0	5.17	1.07	1.05	1.27	0.89
Ln(Patent stock $t - 1$) (tie)	0.56	0.68	0	4.92	0.50	0.68	0.59	0.67

Notes. For *brokerage*, the original measure was transformed using a “2-constraint,” and zero was assigned for inventors without collaborators. For count variables, “1” was added before logging. The final sample used for analysis includes 24,816 inventors. CW, citation weighted.

obtained from the regression. Specifically, from the sample, I generated a random sample of the same size clustered by inventor and estimated the WESML model that includes the *nonselection hazard*. I repeated this process 100 times and computed the standard errors of the estimated coefficients on the explanatory variables.

Given the exponential form of these models, I logged all explanatory variables except for the dummies and the *brokerage* measure. In the tie formation regression, inventors who formed more than one tie appear multiple times. Hence, I clustered the sample by inventor to allow for the errors to be correlated within inventor.

3.4. Analysis 2: Position–Performance Association

In this analysis, I examine how the relationship between brokering position and performance is mediated by inventor-level heterogeneity. Specifically, I compare the coefficients on *brokerage* in relating to performance, estimated from (i) a cross-section model without controls of past performance, (ii) a model with controls of past performance, and (iii) a model with controls of inventor fixed effects. The difference between these coefficients will substantiate the strength of inventor-level heterogeneity in mediating the position–performance association demonstrated in prior studies.

3.4.1. Measures.

Performance Measure. I used *simple patent counts* and *citation-weighted patent counts* as measures of performance. Both measures were adjusted by team size (i.e., number of inventors on each patent).

Brokerage Measure. I used network constraint (Burt 1992) to measure *brokerage*.³⁰ If brokerage facilitates innovative performance, the coefficient on this measure will be positive in the cross-section.

Past Performance. I used *patent stock* for past performance, both simple counts and citation-weighted counts cumulated up to $t - 1$ (team size adjusted). If past performance correlates with *brokerage* through tie formation, controlling for this variable is expected to reveal the mediation effect of performance heterogeneity.

Other Controls. *Number of collaborators* controls for the effect of network size. Collaborating with many inventors may improve performance because the inventor would benefit from access to a broader knowledge base. The *brokerage* measure does not necessarily capture this effect (cf. Fleming et al. 2007). Similarly, *size of inventor pool* controls for differences in the environment for inventive activities. The variable represents technological opportunities that provide a social-capital-like environment (Coleman 1988). It may also capture the competition for innovation opportunities (Fleming et al. 2007). For this variable, I counted the number of inventors in the organizations to which the inventor belonged. *Number of patent classes* controls for differences in the content of network relations. Rodan and Galunic (2004) found that access to heterogeneous knowledge contributes to managerial performance, independent of network position. To proxy for the knowledge heterogeneity of collaboration ties, I counted the number of distinct patent classes the inventor’s patents were assigned to.³¹

Several dummy variables were included to control for differences in the collaboration pattern. Not all ties have the same cost and benefit. Knowledge flows more easily through some types of ties (e.g., Hansen 1999). Organizational boundaries and geographic distance too affect the flow of knowledge (Kogut and Zander 1992, Jaffe et al. 1993). Thus, even if the networks are structurally equivalent, inventors may exhibit different performance depending on the collaboration pattern.

Table 2 Descriptive Statistics (Analysis of Position–Performance Association)

Variables	(N = 116,468)			
	Mean	Std. dev.	Min	Max
Ln(<i>Simple patent count</i>)	0.14	0.28	0	4.65
Ln(<i>Citation-weighted patent count</i>)	0.36	0.79	0	7.31
Ln(<i>Patent stock t – 1</i>)	0.73	0.61	0.05	4.69
Ln(<i>Citation-weighted patent stock t – 1</i>)	2.00	1.33	0	8.57
<i>Brokerage</i>	1.16	0.28	0.06	1.99
Ln(<i>Number of collaborators</i>)	1.39	0.57	0.69	5.17
Ln(<i>Size of inventor pool</i>)	2.72	1.23	0	6.10
Ln(<i>Number of patent classes</i>)	1.34	0.45	0.69	4.56
<i>Coauthorship</i> dummy	0.42	0.49	0	1
<i>Cross-discipline</i> dummy	0.16	0.37	0	1
<i>Cross-region</i> dummy	0.35	0.48	0	1
<i>Cross-country</i> dummy	0.08	0.27	0	1
<i>Cross-organization</i> dummy	0.07	0.25	0	1

Notes. For *brokerage*, the original measure was transformed using a “2-constraint.” For count variables, “1” was added before logging. The final sample used for analysis includes 23,090 inventors.

The *cross-discipline* dummy controls for heterogeneity across scientific disciplines. This dummy indicates whether the inventor has collaborated with nonbiotech inventors. The *cross-region* dummy and *cross-country* dummy control for the effect from geographic distance, and the *cross-organization* dummy controls for the effect of collaborating across organizations. The *coauthorship* dummy controls for the effect from involvement in other forms of knowledge creation. Although there is evidence of conflicting logics between science and innovation (Gittelman and Kogut 2003), many scientific discoveries ultimately get patented (Murray and Stern 2007). Hence, collaboration in other intellectual domains may exert influence on inventor performance. Finally, *year* dummies control for time fixed effects. To avoid simultaneity bias, I lagged all explanatory variables by one year except for year dummies. Table 2 provides descriptive statistics. The final sample included 116,468 inventor-year observations for 23,090 inventors.³²

3.4.2. Estimation Method. In the model, an inventor’s patenting performance is a function of the inventor’s brokerage, inventor-level heterogeneity, and other controls such as time and ego network characteristics. I proxied inventor-level heterogeneity by *patent stock* (random-effects model) and *inventor fixed effects* (fixed-effects model). I used the ordinary least squares (OLS) method to estimate the model, allowing for errors to be correlated within an inventor. Given that the distribution of patent counts—even after adjusted for team size—is skewed, I logged all variables except for the dummies and the *brokerage* measure to reduce the problem of heteroskedasticity. In the results, I report heteroskedasticity-robust standard errors allowing for clustering by inventor.³³

4. Results

4.1. Analysis of Collaboration Tie Formation

4.1.1. Main Results. I begin with the results of how past performance correlates with inventors’ propensities to form certain types of collaboration ties. Table 3 presents the results from the WESML logit analysis. In all models, the dummy for *repeating tie* was omitted—and thus forms the baseline—because not all three type dummies can enter the models due to multicollinearity. Hence, any effect on the explanatory variables involving other types of ties needs to be interpreted with respect to the baseline case. Model (1) only included dyad- and tie-level attributes as controls. The coefficients on both indicators of *redundant tie* and *nonredundant tie* were strongly negative, implying that inventors are much less likely to start a new—either redundant or nonredundant—collaborative relationship than to repeat collaboration with those who they had prior collaborative relationships with.³⁴ A previous coauthorship between the dyad increased the likelihood of patent collaboration. Prior relationships in one form of knowledge creation are thus likely to engender collaboration in other forms of similar activity.³⁵ Coaffiliation and geographical proximity were positive predictors of tie formation. Among the tie-level attributes, new inventors were more likely to receive invitations to collaboration. This also reflects the growth trend of the network as new inventors are added to the population. The likelihood of tie formation increased with the tie’s network size but decreased with the tie’s experience, consistent with Powell et al. (2005).

In Model (2), inventor attributes including past performance were introduced. For the measure of past performance, I used *simple count patent stock*. Note that focal inventor attributes need to be interacted with type dummies for their effects on each type to be specified. This is because the dependent variable is an indicator of forming *any* type of tie, and hence the contingent effect of an attribute on a particular type of tie formation is only identified through an interaction term. This contrasts with other explanatory variables that are dyad or tie specific. Given three types of ties, models can include two interaction terms. The coefficients on the linear term thus represent their impact on the baseline case—i.e., likelihood of repeating a tie. As hypothesized, *patent stock* is positively associated with the probability of forming a nonredundant tie. The coefficient on the interaction between *patent stock* and the *nonredundant tie* dummy was positive and larger than that between *patent stock* and the *redundant tie* dummy. The baseline effect of *patent stock* was negative, indicating that high-performing inventors are less prone to repeat a tie. Thus, inventors with superior performance records are most likely to form a nonredundant collaboration tie. In terms of marginal effect, the coefficient implies that one

Table 3 WESML Logit Regression of Collaboration Tie Formation

Dependent variable = Pr(<i>Tie formation</i>)	(1)	(2)	(3)	(4)	(5)
	Dyad and tie controls	<i>Simple patent stock</i> $t - 1$	<i>CW patent stock</i> $t - 1$	<i>Simple patent stock</i> $t - 1$	<i>CW patent stock</i> $t - 1$
Measure of <i>patent stock</i>					
Dyad attributes					
<i>Redundant tie</i>	-7.285*** (0.085)	-7.698*** (0.245)	-7.851*** (0.248)	-7.699*** (0.245)	-7.852*** (0.248)
<i>Nonredundant tie</i>	-5.687*** (0.090)	-5.993*** (0.175)	-6.196*** (0.179)	-6.286*** (0.233)	-6.520*** (0.237)
<i>Coauthorship dummy</i>	3.494*** (0.137)	3.510*** (0.104)	3.518*** (0.101)	3.500*** (0.107)	3.506*** (0.103)
<i>Coaffiliation dummy</i>	6.024*** (0.041)	6.004*** (0.041)	6.016*** (0.041)	6.001*** (0.041)	6.013*** (0.041)
<i>Colocation dummy</i>	1.217*** (0.052)	1.292*** (0.054)	1.279*** (0.054)	1.298*** (0.054)	1.286*** (0.054)
Tie attributes					
<i>New inventor dummy</i>	0.940*** (0.121)	0.990*** (0.115)	0.995*** (0.113)	0.989*** (0.116)	0.994*** (0.114)
Ln(<i>No. collaborators</i>) (tie)	0.038 (0.069)	0.063 (0.068)	0.076 (0.066)	0.061 (0.069)	0.074 (0.067)
Ln(<i>Patent stock</i> $t - 1$)	-0.270*** (0.039)	-0.213*** (0.044)	-0.225*** (0.043)	-0.210*** (0.044)	-0.222*** (0.043)
Inventor attributes					
<i>Brokerage</i>		0.291 (0.226)	0.305 (0.227)	0.240 (0.228)	0.247 (0.229)
<i>Brokerage</i> × <i>redundant tie</i>		-1.061*** (0.315)	-1.067*** (0.318)	-1.062** (0.316)	-1.070** (0.318)
<i>Brokerage</i> × <i>nonredundant tie</i>		-0.355 (0.275)	-0.413 (0.277)	0.018 (0.326)	-0.054 (0.326)
Ln(<i>No. collaborators</i>)		-0.321** (0.123)	-0.442*** (0.119)	-0.358** (0.126)	-0.484*** (0.122)
Ln(<i>No. collab.</i>) × <i>redundant tie</i>		0.700*** (0.160)	0.797*** (0.155)	0.701*** (0.161)	0.798*** (0.156)
Ln(<i>No. collab.</i>) × <i>nonredundant tie</i>		-0.238 (0.156)	-0.083 (0.152)	-0.329* (0.160)	-0.182 (0.157)
Ln(<i>Patent stock</i>)		-0.395*** (0.071)	-0.148*** (0.036)	-0.423*** (0.073)	-0.163*** (0.037)
Ln(<i>Patent stock</i>) × <i>redundant tie</i>		0.336*** (0.071)	0.130*** (0.036)	0.336*** (0.092)	0.131** (0.047)
Ln(<i>Pat. stock</i>) × <i>nonredundant tie</i>		0.654*** (0.089)	0.282*** (0.045)	0.648*** (0.089)	0.285*** (0.045)
Nonselection hazard (λ)				0.602 (0.342)	0.651 (0.337)
Constant	-5.373*** (0.229)	-4.875*** (0.292)	-4.808*** (0.302)	-4.342*** (0.409)	-4.216*** (0.418)
No. observations	530,460	530,460	530,460	530,460	530,460
Log-likelihood	-935.7	-913.9	-913.7	-913.6	-913.4
McFadden's pseudo- R^2	0.634	0.643	0.643	0.643	0.643

Notes. Bootstrapped standard errors, clustered by inventor, are in parentheses. Year effects are included in all models and are all jointly significant. CW, citation weighted.

* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$. All tests are two-tailed.

unit increase in *patent stock* (in logged value, hereafter) translates into about 65% increase in the likelihood of forming a nonredundant tie.³⁶ In contrast, the negative coefficient on the baseline term suggests that the same increase in *patent stock* is associated with an almost 40% decrease in the probability of repeating a collaboration tie. Model (3) used *citation-weighted patent stock* as the measure of past performance. The results were

analogous to those in Model (2). The magnitude of the coefficient on *patent stock* decreased, but its effect was still sizeable: one unit increase in *patent stock* translated into a 28% higher chance of forming a nonredundant tie and about 15% lower likelihood of repeating a tie.

Among other focal inventor attributes, current brokerage was not a significant correlate of forming a repeating tie and a nonredundant tie, and was negatively associ-

ated with a redundant tie formation. Thus, after controlling for the differences in the dyad and tie attributes, brokering position does not increase the probability of forming a nonredundant tie. Nor does the *number of collaborators* significantly correlate with a nonredundant tie formation, though having more collaborators seems to increase the likelihood of collaborating with a redundant tie.

Models (4) and (5) included the *nonselection hazard* (λ), recovered from the selection equation in the first stage, to control for the differences in patenting hazard across inventors and over time. On both *patent stock* measures, the effect of past performance on collaboration tie formation remained very similar. A unit increase in *patent stock* was associated with a 65% (*simple patent stock*) and a 29% (*citation-weighted patent stock*) higher chance of forming a nonredundant tie, whereas it implied significantly lower likelihoods of repeating a collaboration tie. The coefficient on the *nonselection hazard* was positive but insignificant. Coefficients on other covariates generally stayed the same.

Hence, the results support the hypothesis that inventors with superior track records are more likely to form a collaboration tie that enhances their brokering position. This finding also seems robust to potential selection in collaborative patenting.

4.1.2. Robustness Checks. I ran alternative model specifications varying the proxies for past performance, limiting the sample to a subset, and changing the specification method.³⁷ First, I used depreciated patent stocks as proxies for past performance. Performance in the more distant past may get discounted when inventors evaluate the potential value of a tie. I constructed these measures by cumulating the patents with a 15% depreciation rate (e.g., Cockburn and Griliches 1988). Second, I used patent stocks cumulated only over five years. This limits the window of performance to the previous five years. The results were robust to these variations in past performance proxies. In all specifications, the interaction term between *patent stock* and the indicator of a *nonredundant tie* was significantly positive. A unit increase in the corresponding *patent stock* implied a 28% ~ 75% higher probability of nonredundant tie formation, whereas the same change was associated with a 15% ~ 45% lower likelihoods of repeated collaboration. Hence, the results from the main specifications do not seem to have issues of construct validity.

With the stratified sampling design, the results are unlikely to suffer from the uneven case-control matches. At the time of tie formation, all inventors may not have all alternative types available. As long as the controls correctly represent the risk set associated with each tie formation, the estimates should be unbiased. Nevertheless, for a conservative estimation, I ran the specification limiting the sample to cases with all three types of controls. Even on this narrower sample ($N = 244,060$), past

performance was positively associated with the probability of forming a nonredundant tie. A unit increase in *simple count patent stock* implied a 44% higher probability of forming a nonredundant tie. The positive effect remained significant when the *citation-weighted patent stock* was used, though the magnitude was smaller (15%). Thus, the main results seemed to hold on this restricted sample.

In the logit estimation, I tried to control for important inventor-level covariates of tie formation. However, unobserved heterogeneity in inventor attributes might have affected the coefficient estimates. Thus, I ran conditional logit estimations on the full sample and on the restricted sample, varying the patent stock measures. The conditional logit estimator (McFadden 1974) groups the sample by cases and controls, and computes the likelihood relative to each group. Hence, any effect from the elements common to cases and controls are purged from the estimates. The results were robust to this change in the specification method. On all measures, higher *patent stock* was significantly and positively associated with the probability of forming a nonredundant tie. The model fit also improved. Similar results were obtained from a restricted sample of complete cases and controls. Hence, the omitted variable bias from uncontrolled heterogeneity, if any, does not appear consequential.

Other robustness checks also include (1) using *patent stock* based on unadjusted patent counts (i.e., assign one patent to each collaborating inventor), (2) limiting the sample to inventors with at least one patent assigned to a university (“university inventors”), and (3) limiting the sample to inventors affiliated with large organizations (i.e., organizations with at least 50 inventors in the preceding five-year window). The first one was to examine the effects from giving equal credits to each inventor regardless of team size associated with the patents when constructing *patent stock* measures. The next two robustness checks were to minimize the differences in opportunities or constraints in pursuing collaboration depending on the organizational environments. In all these variations, past performance remained a strongly positive correlate of a nonredundant tie formation.

4.2. Analysis of Position–Performance Association

4.2.1. Main Results. I now turn to estimating the mediation effect of inventor-level heterogeneity on the association between brokering position and patent performance. Table 4 presents the linear regressions of the position–performance association. In Models (6) through (8), *simple patent counts* were used as the dependent variable. Model (6) shows the random-effects coefficients of the variables, without controlling for past performance. Consistent with the findings in prior studies (e.g., Burt 2004, Fleming et al. 2007), the coefficient

Table 4 Linear Regression (OLS) of Position–Performance Association

Dependent variable	Ln(Simple patent count)			Ln(Citation-weighted patent count)		
	(6)	(7)	(8)	(9)	(10)	(11)
<i>Brokerage</i>	0.078*** (0.007)	0.049*** (0.007)	−0.023 (0.012)	0.266*** (0.018)	0.093*** (0.017)	−0.040 (0.032)
Ln(<i>CW patent stock</i> $t - 1$)		0.045*** (0.001)			0.168*** (0.004)	
Ln(<i>No. collaborators</i>)	0.008 (0.004)	0.019*** (0.004)	0.057*** (0.008)	0.003 (0.011)	0.075*** (0.010)	0.038 (0.020)
Ln(<i>Size of inventor pool</i>)	0.001 (0.001)	0.002 (0.001)	0.002 (0.002)	−0.009*** (0.003)	−0.008*** (0.002)	−0.017* (0.007)
Ln(<i>No. patent classes</i>)	0.098*** (0.004)	0.063*** (0.004)	−0.004 (0.006)	0.377*** (0.012)	0.186*** (0.010)	−0.055*** (0.016)
<i>Coauthorship dummy</i>	−0.003 (0.002)	−0.0003 (0.002)	0.006 (0.004)	−0.015* (0.006)	0.001 (0.006)	0.035*** (0.010)
<i>Cross-discipline dummy</i>	−0.012** (0.004)	−0.023*** (0.004)	−0.027*** (0.005)	0.004 (0.011)	−0.053*** (0.012)	−0.082*** (0.013)
<i>Cross-region dummy</i>	0.002 (0.003)	−0.0005 (0.003)	−0.008 (0.005)	0.021** (0.007)	0.006 (0.007)	−0.036* (0.014)
<i>Cross-country dummy</i>	−0.024*** (0.004)	−0.013*** (0.003)	0.008 (0.006)	−0.121*** (0.010)	−0.067*** (0.009)	0.002 (0.016)
<i>Cross-organization dummy</i>	0.011* (0.005)	0.011* (0.005)	0.026** (0.010)	−0.0002 (0.012)	0.002 (0.011)	0.006 (0.024)
Constant	−0.086*** (0.008)	−0.093*** (0.008)	0.137*** (0.012)	−0.337*** (0.023)	−0.310*** (0.022)	0.722*** (0.034)
Inventor fixed effects	No	No	Yes	No	No	Yes
<i>Brokerage</i> coeff. change ^a	—	−0.029	−0.101	—	−0.173	−0.306
<i>p</i> -value ^b	—	0.001	0.000	—	0.000	0.000
No. observations	116,468	116,468	112,677	116,468	116,468	112,677
No. inventors	23,090	23,090	19,299	23,090	23,090	19,299
Adjusted <i>R</i> ²	0.114	0.158	0.319	0.087	0.145	0.277

Notes. Robust standard errors are in parentheses. Year effects are included in all models and are all jointly significant. CW, citation weighted.

^aChange of the coefficient on *brokerage* relative to that from the random-effects model.

^b*p*-value from the test of coefficient difference with the null that the difference is zero.

p* < 0.05; *p* < 0.01; ****p* < 0.001. All tests are two-tailed.

on *brokerage* was significantly positive. Thus, without controls of inventor-level heterogeneity in performance, brokering position appears to correlate positively with patent performance. Knowledge heterogeneity of collaboration ties was positively associated with patent counts, but network size and the pool of potential collaborators were not significant correlate of innovative performance. Among cross-collaboration variables, only cross-organizational collaboration had a positive impact on patenting. Prior coauthorship was not significant in explaining patent performance.

Model (7) included the *citation-weighted patent stock* to control for the effect of past performance on position–performance association.³⁸ The effect was sizeable. With this inclusion, the coefficient on *brokerage* decreased by 37%, and the change was statistically significant (*p*-value = 0.001 on a two-tail test). Coefficients on other controls generally remained similar.

The significant reduction in the coefficient on *patent stock* indicates that inventor-level heterogeneity in

past performance positively mediates the position–performance association. *Patent stock*, however, cannot capture all qualities that vary across inventors and also correlate with both brokering position and patent performance. For instance, some inventors may have gone to schools that provided more opportunities for collaboration and better training. Others may have family backgrounds that promoted exploration and had close connections to great inventive spirits. This multiplicity of inventor-level heterogeneity cannot thus be contained in a few variables. The imperfect reduction of coefficient on *brokerage* suggests that other mediators are also at play (cf. Baron and Kenny 1986).³⁹ Hence, in Model 8, I included *inventor fixed effects* to control for all inventor-level heterogeneity that might mediate the position–performance correlation. With *inventor fixed effects* as controls instead of *patent stock*, the coefficient on *brokerage* fell drastically and became no longer significant at the conventional level (0.049, *p*-value = 0.000 → −0.023, *p*-value = 0.057). Indeed,

inventor-level heterogeneity in underlying quality almost perfectly mediates the positive association between brokering position and patent performance. With the fixed effects, the coefficients on the *number of patent classes* and *cross-country* dummies lost significance, whereas the coefficient on the *cross-discipline* dummy remained negative. Coefficients on other controls were also comparable.

Analogous results were obtained when *citation-weighted patent counts* were used as a performance measure (Models (9) through (11)). In the random-effects model without controls of past performance, the coefficient on *brokerage* was strongly positive (Model (9)). With controls of past performance in Model (10), the effect of *brokerage* on performance decreased considerably and significantly. When inventor fixed effects were controlled for in Model (11), the coefficient on *brokerage* became indistinguishable from zero (-0.04 , p -value = 0.217). With the fixed effects, the coefficient on *brokerage* became strongly positive, whereas the coefficients on *number of collaborators* and the *cross-country* dummy lost significance. Other controls remained qualitatively similar.

From these results, two things seem evident: first, inventor fixed effects are responsible for much of the variations in patent performance; second, inventor-level heterogeneity in underlying quality positively mediates the position–performance relationship in the collaborative patenting network. Accounting for the differences in the effect of inventor quality on brokering position and performance thus rendered the association between brokerage and performance almost nonexistent. These support Hypothesis 2.

4.2.2. Robustness Checks. I performed several robustness checks, varying the measures of brokerage and performance.⁴⁰ First, I used *efficiency* and *density* as alternative measures of *brokerage*.⁴¹ These are often used in the literature as proxies for *brokerage* or an inverse thereof (Burt 1992, Podolny and Baron 1997, Nerkar and Paruchuri 2005, Obstfeld 2005), albeit less widely than *network constraint*. On both of these alternative measures, the changes in the coefficients on *brokerage* due to controls of past performance were sizable and significant. Moreover, with controls for *inventor fixed effects*, the coefficients on brokerage dropped further and, in most cases, became insignificant.⁴² Coefficients on other controls remained qualitatively similar.

In the main specifications, I controlled for either time-varying effects of inventor-level heterogeneity—via *patent stock*—or time-invariant effects thereof—via *inventor fixed effects*—but not both. However, it is possible that innovation performance owes in part to inventor-level capability that changes over time as well as to time-invariant characteristics of inventors. For instance,

inventors may learn over time to become better at innovating, which is largely reflected in *patent stock*. Then, one may want to control for both of these effects in the estimation. To incorporate these possible dynamics in the performance–position feedback while accounting for the effects of inventor-fixed heterogeneity, I ran a dynamic panel data model with fixed effects (Arellano and Bond 1991).⁴³ For both measures of performance, the changes in the coefficients on *brokerage* were more drastic when lagged patent counts and inventor fixed effects were simultaneously included as controls. The coefficients on *brokerage*, which were significantly positive without those controls, became insignificant for the *citation-weighted patent counts* (-0.071 , p -value = 0.476), and even turned negative for the *simple patent counts* (-0.187 , p -value = 0.000). The coefficients on lagged *patent counts* were all positive and significant, indicating that performance excellence tends to persist.

Thus far, I have used team-size-adjusted patent counts to proxy for performance. An alternative construction of the measure would credit one patent to each collaborating inventor. Then, the dependent variables become nonnegative integer values. Moreover, the distribution of these patent counts exhibits overdispersion. The negative binomial model on count data generalizes the Poisson model to allow for such overdispersion and estimates the coefficients consistently (Hausman et al. 1984). Using this model, I estimated random-effects coefficients and fixed-effects coefficients, with the latter conditional upon the sum of the dependent variable. The results were comparable to those in the main specifications. In cross-section, the coefficients on *brokerage* were significantly positive. With past performance control, these coefficients dropped considerably, though both stayed significant. When the *inventor fixed effects* were controlled for, the coefficients on *brokerage* either reduced to essentially zero, for the *citation-weighted patent counts* (0.039 , p -value = 0.386), or turned negative, for the *simple patent counts* (-0.089 , p -value = 0.047).⁴⁴

Other robustness checks also include (1) varying the time window of network relationships (three-year window and seven-year window); (2) using brokerage measures based on dichotomized data; (3) using citation-weighted patent counts excluding self-cites; (4) dropping all observations with only one collaborator; (5) dropping inventors with all zero outcomes;⁴⁵ and (6) randomly selecting a subset of inventors.⁴⁶ In all of these alternative specifications, the main results held robustly.

5. Discussion

The analysis of tie formation by biotech inventors showed that high-performing actors are more apt to form collaboration ties that lead them to occupy brokering positions. Past performance was significantly and positively associated with the probability that inventors

form a nonredundant collaboration tie: a unit increase in patent stock led to 30% ~ 65% higher chances of forming a nonredundant tie. By demonstrating that actor-level attributes—performance history in particular—strongly predict the pattern of future tie formation, this analysis substantiates a process by which actors arrive at certain network positions. This process has been largely neglected in the literature. Without controlling for the feedback from past performance, brokering position was positively correlated with patent performance. However, when the performance heterogeneity across inventors was accounted for, the positive position–performance association disappeared and, in some specifications, even turned sign. These results were robust to variations in measures and specification methods. The analysis reveals that, in the collaborative inventor network, actor-level heterogeneity strongly mediates the observed association between brokerage and performance. Hence, controlling for this heterogeneity appears critical when explicating the relationship between the two. In this study, the use of large-scale longitudinal data on patent collaboration facilitated such controls, thereby minimizing the potentially confounding effects from unobserved factors, particularly those from actor-level performance heterogeneity.

The findings of this study point to a potentially important issue in social networks research: endogeneity in network position. Though this study only concerned brokerage, its fundamental idea may well extend to other characterizations of actor position. To the extent that network positions arise from actor-level attributes that also correlate with the economic outcomes of actors, care must be taken in interpreting the observed relationship between the position and actor performance. It may be that a certain network position—such as brokerage—appears valuable because it is available only to the actors who outperform others. This study provides some evidence to this plausible account, for which relatively few empirical endeavors have been made in the literature.

In this exercise, I have focused on one possible source of endogeneity, i.e., the feedback from past performance. Another possible source of endogeneity is the strategic intentions of actors in forming network ties. When deciding on exchange relationships, actors consider the possible consequences of such decisions on their economic performance. In fact, most theoretical models on network formation assume that actors strategically form or sever ties to maximize the final payoffs (e.g., Bala and Goyal 2000, Ryall and Sorenson 2007). Under this scenario, the resulting network positions will be endogenous to performance because actors optimally choose positions with respect to their payoffs. Despite the difference in their mechanisms, however, the two sources of endogeneity are not necessarily exclusive of each other. Whether actors differ in their ability to behave strategically or they face different constraints in forming ties,

actor-level heterogeneity does play a role in shaping the observed position–performance relationship. Hence, the empirical tests in this study have, at least partially, incorporated the repercussions from the other source of endogeneity.

A few potential limitations merit discussion. First, I treated the collaborative networks among U.S. inventors as one macronetwork. There are many subfields within biotechnology that exhibit technological dissimilarities to each other. This may cause fragmentation of the network, limiting extensive brokerage. However, as long as network position is defined relatively—and these subnetworks are more similar to each other than to those outside the discipline—treating them as a macronetwork seems reasonable. Second, the network relationships considered in this study are not complete. Other socioeconomic factors could also influence individuals' inventive activity. Nevertheless, this study captures the two most relevant relationships of inventors: copatenting and coauthorship. Thus, the findings are less likely to suffer from the bias due to neglecting other relationships.

6. Conclusion

This study demonstrates that, in the context of brokerage, actor-level attributes such as performance history largely drive the positional asymmetry across actors. It is the high-performing actors who come to occupy brokering positions because they enjoy advantages in forming network ties that lead to such positions. The findings of this study thus suggest that the causal relationship between the two is essentially endogenous. At least in the collaborative patenting among biotech inventors, the association between brokerage and innovative performance appears attributable almost entirely to the mediating actor-level heterogeneity. This performance–position feedback may have—at least partially—confounded the effects of brokering position on subsequent performance that were often demonstrated in prior studies. Absent the mediation, one may not observe significant performance advantages arising from occupying brokering position.

This study relates to the networks literature in several ways. First of all, the finding that actor-level attributes interact with network position complements that of earlier qualitative studies that document the existence of key individuals performing boundary roles in interorganizational communication networks (Allen 1977, Tushman 1977). Thus, this study reinforces the strategic importance of identifying and promoting these individuals when organizing the process of innovation, particularly for the organizations aiming to maximize the efficacy of their intra- and interorganizational networks.

The study also speaks to the practice of value-enhancing strategy through network formation. Likely motivated by the findings in the literature, pursuing a

career enhancement or an improvement of organizational effectiveness through network building is becoming a popular strategy. Despite the advertised value of such strategy, this study predicts that the returns on such investment most likely accrue to those who begin with better capabilities. The finding thus resonates with prior studies such as that by Ryall and Sorenson (2007), who show that positional advantage not accompanied by underlying resources is at most ephemeral, and Zaheer and Bell (2005), who demonstrate that the benefits of network position are contingent upon underlying capabilities. Hence, one may need to evaluate the validity of such network-building strategies more carefully before putting them into practice.

The findings notwithstanding, this study does not deny the forces that network relationships exert on the behavior of actors. Social networks in general undoubtedly influence actors' decisions and behavior, because they affect the costs and benefits actors face when deciding which set of actions to exercise (Jackson 2006). Differences in the associated cost and benefit will then lead to differentials in the final payoffs. In this sense, network structures affect performance. However, to maximize the learning from the conclusions based on observed relationships, one must begin by clearly understanding the mechanisms that generate positional asymmetry among actors. It is on this point that this study hopes to have made a contribution to the literature.

Acknowledgments

The author thanks Olav Sorenson, John de Figueiredo, Marvin Lieberman, Michael Darby, and Phil Bonacich, Lee Fleming, and two anonymous reviewers for their valuable advice and suggestions. Thanks are also due to Manuel Trajtenberg for generously sharing the inventor data. Insightful comments from Natarajan Balasubramanian, Min Ha Hwang, Gabriel Natividad, and the seminar participants at University of California, Los Angeles Anderson School, Georgia Institute of Technology, and 2007 Academy of Management Meeting are appreciated. The usual disclaimer applies.

Appendix. Inventor Identification Procedure

Because inventors are the basis of collaboration networks in this study, it is crucial to correctly identify individual inventors and match them to patents. Trajtenberg et al. (2006), from whom I obtained initial inventor identifications, developed an algorithm to match all inventors in the NBER patent data set (Hall et al. 2001) to over 2.1 million patents, and identified 1.6 million unique inventors associated with these patents. This name-matching algorithm applies a number of scoring criteria based on an inventor's middle name, geographic location (cities in particular), technological area (i.e., patent class), assignee, identity of coinventors, and relative frequency of the name to the standardized inventor names pregrouped by the Soundex system (which addresses variations in the spelling). Though a pioneering work, however, this algorithm is not error free. In particular, the algorithm performs very poorly for inventors with non-English based

names. Thus, I ran visual inspections on each inventor record and made corrections when misidentification was apparent. The relatively small size of the sample (45,537 inventors—including nonbiotech inventors—on 57,605 patents in total) made these manual adjustments containable. For this task, I relied on three sources of information: patent documents, personal information services, and journal publication records. For each inventor, I first checked from the USPTO website (www.uspto.gov) whether the person appears with the same set of coinventors and the same assignee. If so, I treated the inventor as the same person even if the inventor appears under slightly different names (most likely due to misspelling). For inventors with only one patent or with no coinventors, this method could not be used. Thus, for these cases, I searched the inventor name on USA People Search (www.usa-people-search.com), an online information services firm that provides access to personal data collected from various public records such as taxes, utilities, marriage, etc. The free-version search returns basic information such as name (along with a.k.a.'s if available), age, city, and state. I cross-checked these data with those on the patent document to determine whether the inventor was correctly identified, and made corrections as needed. This method was not valid for inventors who have never lived in the United States. The use of journal publication records addressed this problem for the most part. Using the combination of the last name and the initials of first and middle names, I searched the ISI citation database to identify all scientific journal publications associated with the inventors in the sample (these data were also used for constructing control variables in the analysis). Each article in this database provides the name of the correspondence author and the person's contact information including address and organizational affiliation. I cross-checked this information with that on the patent document to determine whether the inventor was correctly identified and made corrections as needed. Though these corrective steps might not guarantee perfect identifications, I am confident that any potential biases from misidentification were kept to a minimum.

Endnotes

¹See Jackson (2006) for a review of recent developments in this literature.

²From a semantic standpoint, quality can imply anything. I use quality mostly as a synonym for capability.

³Actors also choose between collaboration and independent work to perform a given task, and actor-level quality may systematically correlate with, and in an extreme may even completely determine the mode of performing the task—e.g., high-performing actors only choose to work independently, whereas low performers rely solely on collaboration, or vice versa. Even in this case, however, the argument should still hold for the actors that choose to engage in collaboration. That is, assuming sufficient heterogeneity in quality among the actors that collaborate on a task, actors with relatively higher quality should form ties with greater returns in order to compensate for the higher opportunity cost they have to incur for such collaboration. Actors that only choose to work independently are outside the focus of this study and hence drop out of the equation.

⁴Here, I focus on the positive outcome of brokerage that derives from the strategic value of nonredundant ties. Research

also indicates that, in certain contexts, cohesion characterized by many strong, redundant ties proves more beneficial (e.g., Obstfeld 2005, Fleming et al. 2007).

⁵The initial performance triggering the amplification may not always come from superior quality. A random “shock” in performance uncorrelated with underlying quality can also generate a persistent effect. In an extreme, even a fraud can trigger the cycle, as the following anecdote testifies. In early 2001, a young physicist in Bell Labs, named Jan Hendrik Schön, announced in *Nature* that he had succeeded in producing an organism-based transistor on the molecular scale. His work was such a ground-breaking achievement that he quickly rose to prominence and was regarded as a promising candidate for a Nobel Prize. A herd of physicists soon became his coauthors; in 2001 only, he produced one paper every eight days, collaborating with more than 20 coauthors. This continued in 2002, when Bell Labs began an investigation of possible fraud. The investigation revealed that none of the results reported in his papers was verifiable. Bell Labs fired him immediately and most of his papers once regarded as breakthroughs were withdrawn. Adding to his shame, his alma mater canceled his Ph.D. degree in 2004.

⁶The original data set contained patents granted through 1999. Hall extended it to include patents through 2002 (<http://elsa.berkeley.edu/~bhhall/bhdata.html>). An update through 2006 was recently completed (<https://sites.google.com/site/patentdataprotect/home>).

⁷Considerable time lags often exist between scientific discovery and patent filing in the biotech industry (Murray and Stern 2007). For instance, it was only in 1979 that the 1973 discovery of recombinant DNA was filed for a patent.

⁸For the 3.4 million patents granted during 1963–2002, the average time lag between application and grant is 2.1 years (median, 2 years; standard deviation, 1.1 years).

⁹I thank Manuel Trajtenberg for sharing these initial data with inventor identification.

¹⁰There are 17,270 inventors associated with these nonbiotech patents. These nonbiotech inventors were only used for computing network measures of the biotech inventors.

¹¹An N -way collaboration (i.e., a patent with N inventors) was decomposed into $N(N - 1)/2$ dyadic relationships. Inventors with no patent collaboration throughout their careers were excluded from the sample.

¹²Of 200,515 cases, 62,413 (31%) had three controls, 44,016 (22%) two controls, and 94,086 (47%) one control.

¹³This is true when the computation of brokerage is based on dichotomized network data. With valued network data, the measure can increase even with the formation of a repeated tie, albeit minimally.

¹⁴The effect of closing ties on brokerage is largely indeterminate. The direction of change depends on the distribution of prior links between the closing tie and the focal inventor’s collaborators. If the closing tie was previously linked to many of the collaborators, it may decrease the focal inventor’s brokerage. But if only a few prior links existed, the focal inventor’s brokerage may increase with the closing tie formation. Both scenarios depend also on the measure used for quantification.

¹⁵One would ideally like to assign different weights to each inventor based on individual contribution. However, the patent document provides no information in this respect. Hence, I treated all inventors as equal contributors.

¹⁶I also constructed *patent stock* with a depreciation of 15% per annum and *patent stock* cumulated over only five years prior to t . These measures were used for robustness checks.

¹⁷Burt et al. (1998, p. 82) claim that “network constraint has the construct validity of accumulated evidence showing the expected (negative) association between constraint and (manager) performance” (parentheses added).

¹⁸The proportion $p_{ij} = z_{ij} / \sum_q z_{iq}$, where z_{ij} measures the intensity of i ’s relationship with j . On valued network data, z_{ij} is the number of collaborations between i and j . On dichotomized network data, z_{ij} is an indicator of collaboration between i and j . I used valued network data for analysis. The two methods, however, yield numerical values that are very similar (correlation coefficient: 0.94) and the empirical results were robust to the choice of computation method.

¹⁹In the sample, the highest value of *network constraint* was 1.936. Also, the measure always takes a positive number. The transformed variable is bounded from above (2) and below (0).

²⁰In the sample, the *coauthorship* dummy was positively correlated with the probability of repeating a tie, whereas it was negatively correlated with the probability of forming a nonredundant tie, though the correlations were moderate (0.207 and -0.200 , respectively).

²¹See Lee (2007, pp. 143–144) for details of this identification process.

²²Murray and Stern (2007) report that, in their sample, there was an average 37.5-month gap between scientific publication and patent grant. Thus, I began from 1974 for identifying coauthorship.

²³For the single-assignee patents, I assumed that all inventors belong to the same organization. Among 57,605 patents in the sample, 2,096 (3.6%) patents have more than one assignee.

²⁴To define geographic regions, I followed the National Science Foundation’s (NSF) definition of regions. The NSF divides the U.S. territory into nine regions, identical to the regional classification used in the census (National Science Foundation 2006). To determine inventor locations, I used the inventor geography information on patent documents (city, state, or country).

²⁵Inventors appearing for the first time in the sample have missing values for some variables. I assigned zero for *patent stock* and network-based measures for each inventor’s first year. Others (6.9%) were dropped from the sample.

²⁶The underlying model for specification takes the following form:

$$Y_{ijt}^* = \beta' X_{ijt} + \gamma' W_{jt} + \alpha' Z_{it} + \varepsilon_{ijt},$$

where Y_{ijt}^* = indirect utility for inventor i from choosing inventor j at time t ; $Y_{ijt} = 1$ if i chooses j at time t ; $Y_{ijt} = 0$ otherwise; $X_{ijt} = i - j$ dyad attributes at time t ; $W_{jt} =$ inventor j ’s characteristics at time t ; $Z_{it} =$ inventor i ’s attributes at time t .

Then, the probability that inventor i chooses inventor j for collaboration becomes

$$Pr(Y_{ijt} = 1 | X_{ijt}, Z_{it}, W_{jt}) = \frac{e^{\beta' X_{ijt} + \gamma' W_{jt} + \alpha' Z_{it}}}{1 + e^{\beta' X_{ijt} + \gamma' W_{jt} + \alpha' Z_{it}}}.$$

The coefficients are estimated by maximizing the following likelihood function:

$$\ln L = \sum_{j \in S} w_j \ln F(\lambda_{ijt}) + \sum_{j \notin S} w_j \ln \{1 - F(\lambda_{ijt})\},$$

where S is the set of j such that $y_j = 1$, and w_j is the inverse of the probability that an observation j is sampled, $\lambda_{ijt} = \beta' X_{ijt} + \gamma' W_{jt} + \alpha' Z_{it}$ and $F(z) = e^z / (1 + e^z)$.

²⁷I thank Lee Fleming and two anonymous reviewers for pointing this out.

²⁸Due to space constraints, the estimated selection equation is not reported here but is available from the author.

²⁹By this, I essentially treat “no patent” and “patent alone” equivalently because neither leads to any tie formation.

³⁰In my sample, there is a significant correlation between the number of nonredundant tie formations and the network constraint measure at the inventor level ($\rho = 0.664$, p -value = 0.000). Thus, it seems valid to use network constraint as a measure of brokering position in Analysis 2, whereas, in Analysis 1, brokerage is proxied by the tendency of forming nonredundant collaboration ties.

³¹A patent may be assigned to multiple patent classes. However, the NBER patent data set contains only the first class. Thus, I identified all other classes of the sample patents from the USPTO website (www.uspto.gov).

³²In total, 4,952 inventors were dropped from the sample due to missing values. For network position to be measured, an inventor should have at least one collaborator. Thus, all isolates were dropped. Inventors who patented first in 1995—which is the end of the sample period—do not have position measure and hence were dropped.

³³For a robustness check, I ran the negative binomial specification on the same model except that patent counts were rounded to the nearest integers. The results were robust to this alternative specification.

³⁴To be precise, these coefficients indicate the likelihood that, conditioning on a tie occurring, it is a redundant tie or a nonredundant tie. However, for a more intuitive understanding, I describe the coefficients as if they represent the probabilities that inventors *choose* certain types of ties.

³⁵To the extent that article coauthorship correlates with patent collaboration, including the *coauthorship* dummy as an explanatory variable may bias the estimates due to the violation of strict exogeneity assumption. To address this possibility, I estimated the models excluding the *coauthorship* dummy. The results were robust to this exclusion, indicating that the correlation is unlikely to be an issue.

³⁶The marginal effects were computed using the “mfx” command in Stata 10 after each estimation. Not all marginal effects are reported here due to space considerations. However, in the case of rare event logit such as here, the marginal effects can be simply approximated by the coefficient estimates. To see this, let $p = \text{prob}[Y = 1]$. Because $\log \text{it}(p) = \log(p/(1-p)) = x_i' \beta$, where x_i' is a vector of regressors and β is a vector of coefficients to be estimated, $p = \exp(x_i' \beta) / (1 + \exp(x_i' \beta))$. It can be shown that the marginal effect of a variable x_i is then $\partial p / \partial x_i = \beta_i p (1 - p)$. For rare events where p is sufficiently small, the right-hand side of the equation is close enough to $\beta_i p$. Therefore, the relative magnitude of the marginal effect with respect to the reference probability (e.g., the mean predicted probability) is approximately $\beta_i p / p = \beta_i$, which is equivalent to the estimated coefficient on x_i .

³⁷Due to space constraints, these results are not reported here but are available from the author.

³⁸Using *simple patent stock* as the measure of past performance produced very similar results.

³⁹Including *patent stock* as a control may also introduce bias due to autocorrelation, although the direction of bias is not certain. Instrumenting for past performance would at least help avoid this particular issue. However, I was not able to find a proper instrument.

⁴⁰Due to space constraints, these results are not reported here but are available from the author.

⁴¹*Efficiency* captures the nonredundancy of an actor’s network ties and is measured as the number of nonredundant ties divided by the total number of ties. *Density* captures the cohesiveness of an actor’s ego network and is measured as the number of links among the ties divided by the number of all possible links among them (links to the focal actor are excluded). Because *density* represents an opposite concept of *brokerage*, I transformed *density* by subtracting it from 1 for a more intuitive interpretation.

⁴²The coefficient on *efficiency* stayed significant even with controls of fixed effects when the *simple patent count* was used as the performance measure. However, the coefficient reduction was still substantial (0.432 \rightarrow 0.128).

⁴³In the presence of a serial correlation, the coefficients with the lagged dependent variable as a regressor are biased. The Arellano-Bond Generalized Method of Moments (GMM) estimator uses lagged values of the dependent variable and other covariates as instruments, thereby producing consistent estimates. I treated all explanatory variables (except for year dummies) as predetermined, and used five years for lag structure to ensure that the model incorporates sufficient dynamics.

⁴⁴Note that these coefficients are not directly comparable to those in the main specifications because they come from a different estimation method. In terms of the instance rate ratio, they correspond to 1.04 and 0.915, respectively.

⁴⁵For both performance measures, the coefficient on brokerage turned insignificant with fixed-effects controls (*simple patent counts*, 0.077*** \rightarrow -0.23, $p = 0.061$; *citation-weighted patent counts*, 0.257*** \rightarrow -0.039, $p = 0.22$).

⁴⁶Using a population of inventors may artificially create network autocorrelation because some of the variables are the same for the inventors collaborating on a patent. In particular, the bipartite structure in this kind of network (cf. Uzzi and Spiro 2005) might lead to a biased result to the extent that the sample overrepresents certain network positions by repeating them. Hence, I ran the same analysis on random samples of inventors, using 20% and 33% as the sampling proportion (the average degree of inventors during the study period ranged from 2.5 to 3.5). The results on these subsamples were similar to those from the main analysis.

References

- Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Admin. Sci. Quart.* **45**(3) 425–455.
- Allen, T. J. 1977. *Managing the Flow of Technology: Technology Transfer and the Dissemination of Technological Information within the R&D Organization*. Massachusetts Institute of Technology, Cambridge.
- Arellano, M., S. Bond. 1991. Some tests of specification for panel data Monte Carlo evidence and an application to employment equations. *Rev. Econom. Stud.* **58**(2) 277–297.
- Bala, V., S. Goyal. 2000. A noncooperative model of network formation. *Econometrica* **68**(5) 1181–1229.

- Baron, R. M., D. A. Kenny. 1986. The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *J. Personality Soc. Psych.* **51** 1173–1182.
- Baum, J. A. C., T. J. Rowley, A. V. Shipilov, Y. Chuang. 2005. Dancing with strangers: Aspiration performance and the search for underwriting syndicate partners. *Admin. Sci. Quart.* **50**(4) 536–575.
- Burt, R. S. 1992. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge, MA.
- Burt, R. S. 1997. The contingent value of social capital. *Admin. Sci. Quart.* **42**(2) 339–365.
- Burt, R. S. 2000. Decay functions. *Soc. Networks* **22** 1–28.
- Burt, R. S. 2004. Structural holes and good ideas. *Amer. J. Sociol.* **110**(2) 349–399.
- Burt, R. S., J. E. Jannotta, J. T. Mahoney. 1998. Personality correlates of structural holes. *Soc. Networks* **20** 63–87.
- Chan, S. H., J. W. Kensinger, A. J. Keown, J. D. Martin. 1997. Do strategic alliances create value? *J. Financial Econom.* **46** 199–221.
- Cockburn, I., Z. Griliches. 1988. Industry effects and appropriability measures in the stock market's valuation of R&D and patents. *Amer. Econom. Rev.* **78**(2) 419–423.
- Cohen, W. M., D. A. Levinthal. 1990. Absorptive capacity: A new perspective on learning and innovation. *Admin. Sci. Quart.* **35**(1) 128–152.
- Cohen, W. M., R. R. Nelson, J. P. Walsh. 2000. Protecting their intellectual assets: Appropriability conditions and why U.S. manufacturing firms patent (or not). Working Paper 7552, National Bureau of Economic Research, Cambridge, MA.
- Coleman, J. S. 1988. Social capital in the creation of human capital. *Amer. J. Sociol.* **94** S95–S120.
- Cosslett, S. R. 1981. Maximum likelihood estimator for choice-based samples. *Econometrica* **49**(5) 1289–1316.
- Fernandez, R. M., E. J. Castilla, P. Moore. 2000. Social capital at work: Networks and employment at a phone center. *Amer. J. Sociol.* **105**(5) 1288–1356.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47**(1) 117–132.
- Fleming, L., S. Mingo, D. Chen. 2007. Brokerage vs. cohesion and collaborative creativity: An evolutionary revolution. *Admin. Sci. Quart.* **52** 443–475.
- Gittelman, M., B. Kogut. 2003. Does good science lead to valuable knowledge? Biotechnology firms and the evolutionary logic of citation patterns. *Management Sci.* **49**(4) 366–382.
- Gompers, P., A. Kovner, J. Lerner, D. Scharfstein. 2006. Skill vs. luck in entrepreneurship and venture capital: Evidence from serial entrepreneurs. Working Paper 12592, National Bureau of Economic Research, Cambridge, MA.
- Gould, R. V. 2002. The origin of status hierarchies: A formal theory and empirical test. *Amer. J. Sociol.* **107**(5) 1143–1178.
- Griliches, Z. 1990. Patent statistics as economic indicators: A survey. *J. Econom. Literature* **27** 1661–1707.
- Gulati, R., M. Gargiulo. 1999. Where do interorganizational networks come from? *Amer. J. Sociol.* **104**(5) 1439–1493.
- Hall, B. H., A. Jaffe, M. Trajtenberg. 2001. The NBER patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research, Cambridge, MA.
- Hall, B. H., A. Jaffe, M. Trajtenberg. 2005. Market value and patent citations. *RAND J. Econom.* **36**(1) 16–38.
- Hamilton, B., J. Nickerson. 2003. Correcting for endogeneity in strategic management research. *Strategic Organ.* **1** 51–78.
- Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Admin. Sci. Quart.* **44** 82–111.
- Hargadon, A., R. I. Sutton. 1997. Technology brokering and innovation in a product development firm. *Admin. Sci. Quart.* **42** 716–749.
- Hausman, J., B. H. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica* **52**(4) 909–938.
- Heckman, J. J. 1979. Sample selection bias as a specification error. *Econometrica* **47** 153–161.
- Jackson, M. O. 2006. The economics of social networks. R. Blundell, W. Newey, T. Persson, eds. *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress, Vol. 1. Cambridge University Press, Cambridge, MA, 1–56.
- Jaffe, A. B., M. Trajtenberg, R. Henderson. 1993. Geographic localization of knowledge spillovers, as evidenced by patent citations. *Quart. J. Econom.* **108**(3) 577–598.
- Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* **3**(3) 383–397.
- Lanjouw, J. O., M. Schankerman. 2004. Patent quality and research productivity: Measuring innovation with multiple indicators. *Econom. J.* **114** 441–465.
- Lee, J. 2007. Essays on the economics of exchange networks. Unpublished doctoral dissertation. University of California, Los Angeles.
- Lotka, A. J. 1926. The frequency distribution of scientific productivity. *J. Washington Acad. Sci.* **16** 317–324.
- Manski, C. F., S. R. Lerman. 1977. The estimation of choice probabilities from choice based samples. *Econometrica* **45**(8) 1977–1988.
- McFadden, D. 1974. Conditional logit analysis of qualitative choice behavior. P. Zarembka, ed. *Frontiers in Econometrics*. Free Press, New York, 105–142.
- Merton, R. K. 1968. The Matthew effect in science. *Science* **159** 56–63.
- Merton, R. K. 1988. The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *Isis* **79** 606–623.
- Murray, F., S. Stern. 2007. Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. *J. Econom. Behav. Org.* **63**(4) 648–687.
- National Science Foundation. 2006. Academic Research and Development Expenditures: Fiscal Year 2004. Report, Division of Science Resources Statistics, National Science Foundation, Arlington, VA.
- Nerkar A., S. Paruchuri. 2005. Evolution of R&D capabilities: The role of knowledge networks within a firm. *Management. Sci.* **51**(5) 771–785.
- Obstfeld, D. 2005. Social networks, the *Tertius lungen* orientation, and involvement in innovation. *Admin. Sci. Quart.* **50** 100–130.
- Podolny, J. M., J. N. Baron. 1997. Resources and relationships: Social networks and mobility in the workplace. *Amer. Sociol. Rev.* **62**(5) 673–693.

- Powell, W. W., K. W. Koput, L. Smith-Doerr. 1996. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Admin. Sci. Quart.* **41**(1) 116–145.
- Powell, W. W., D. R. White, K. W. Koput, J. Owen-Smith. 2005. Network dynamics and field evolution: The growth of interorganizational collaboration in the life sciences. *Amer. J. Sociol.* **110**(4) 1132–1205.
- Reagans, R., B. McEvily. 2003. Network structure and knowledge transfer: The effects of cohesion and range. *Admin. Sci. Quart.* **48** 240–267.
- Reagans, R., E. W. Zuckerman. 2001. Networks, diversity, and productivity: The social capital of corporate R&D teams. *Organ. Sci.* **12**(4) 502–517.
- Rodan, S., C. Galunic. 2004. More than network structure: How knowledge heterogeneity influences managerial performance and innovativeness. *Strategic Management J.* **25**(6) 541–562.
- Rosenkopf, L., A. Nerkar. 2001. Beyond local search: Boundary-spanning exploration, and impact in the optical disk industry. *Strategic Management J.* **22** 287–306.
- Ryall, M. D., O. Sorenson. 2007. Brokers and competitive advantage. *Management Sci.* **53**(4) 566–583.
- Shipilov, A.V. 2006. Network strategies and performance of Canadian investment banks. *Acad. Management J.* **49**(3) 590–604.
- Singh, J. 2005. Collaborative networks as determinants of knowledge diffusion patterns. *Management Sci.* **51**(5) 756–770.
- Soda, G., A. Usai, A. Zaheer. 2004. Network memory: The influence of past and current networks on performance. *Acad. Management J.* **47**(6) 893–906.
- Sorenson, O., T. Stuart. 2001. Syndication networks and the spatial distribution of venture capital investments. *Amer. J. Sociol.* **106**(6) 1546–1588.
- Trajtenberg, M. 1990. A penny for your quotes: Patent citations and the value of innovations. *RAND J. Econom.* **21**(1) 172–187.
- Trajtenberg, M., G. Shiff, R. Melamed. 2006. The names game: Harnessing inventors' patent data for economic research. Working Paper 12479, National Bureau of Economic Research, Cambridge, MA.
- Tushman, M. L. 1977. Special boundary roles in the innovation process. *Admin. Sci. Quart.* **22**(4) 587–605.
- United States Patent and Trademark Office. 2002. Technology Profile Report: Patent Examining Technology Center Groups 1630–1650, Biotechnology. Patent Technology Monitoring Division Report, United States Patent and Trademark Office, Alexandria, VA.
- Uzzi, B., J. Spiro. 2005. Collaboration and creativity: The small world problem. *Amer. J. Sociol.* **111**(2) 447–504.
- Zaheer, A., G. G. Bell. 2005. Benefiting from network position: Firm capabilities, structural holes, and performance. *Strategic Management J.* **26** 809–825.
- Zucker, L. G., M. R. Darby. 2006. Movement of star scientists and engineers and high-tech firm entry. Working Paper 12172, National Bureau of Economic Research, Cambridge, MA.
- Zucker, L. G., M. R. Darby, M. B. Brewer. 1998. Intellectual human capital and the birth of U.S. biotechnology enterprise. *Amer. Econom. Rev.* **88** 290–306.
- Zuckerman, H. 1967. Nobel laureates in science: Patterns of productivity, collaboration, and authorship. *Amer. Sociol. Rev.* **23**(3) 391–403.