Open Innovation and Strategy

Henry W. Chesbrough
Melissa M. Appleyard

A new breed of innovation—open innovation—is forcing firms to reassess their leadership positions, which reflect the performance outcomes of their business strategies. It is timely to juxtapose some new phenomena in innovation with the traditional academic view of business strategy. More specifically, we wish to examine the increasing adoption of more open approaches to innovation, and see how well this adoption can be explained with theories of business strategy. In our view, open innovation is creating new empirical phenomena that exist uneasily with well-established theories of business strategy. Traditional business strategy has guided firms to develop defensible positions against the forces of competition and power in the value chain, implying the importance of constructing barriers to competition, rather than promoting openness. Recently, however, firms and even whole industries, such as the software industry, are experimenting with novel business models based on harnessing collective creativity through open innovation. The apparent success of some of these experiments challenges prevailing views of strategy.

At the same time, recent developments indicate that many of these experimenters now are grappling with issues related to value capture and sustainability of their business models, as well as issues of corporate influence and the potential co-option of open initiatives. In our view, the implications of these issues bring us back to traditional business strategy, which can inform the quest

Chesbrough received support from the Center for Open Innovation at the Haas School of Business and the Alfred P. Sloan Foundation. Appleyard received support from the National Science Foundation under Grant No. 0438736. Jon Perr and Patrick Sullivan ably assisted with the interviews of Open Source Software leaders. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the above funding sources or any other individuals or organizations.

for sustainable business models. If we are to make strategic sense of innovation communities, ecosystems, networks, and their implications for competitive advantage, we need a new approach to strategy—what we call "open strategy."

Open strategy balances the tenets of traditional business strategy with the promise of open innovation. It embraces the benefits of openness as a means of expanding value creation for organizations. It places certain limits on traditional business models when those limits are necessary to foster greater adoption of an innovation approach. Open strategy also introduces new business models based on invention and coordination undertaken within a community of innovators. At the same time, though, open strategy is realistic about the need to sustain open innovation approaches over time. Sustaining a business model requires a means to capture a portion of the value created from innovation. Effective open strategy will balance value capture and value creation, instead of losing sight of value capture during the pursuit of innovation. Open strategy is an important approach for those who wish to lead through innovation.

The Insights and Limits of Traditional Business Strategy

Business strategy is a wide and diverse field. The origins of the concept hearken back to Alfred Chandler's seminal *Strategy and Structure*, where he presented the first systematic and comparative account of growth and change in the modern industrial corporation.[1] He showed how the challenges of diversity implicit in a strategy of growth called for imaginative responses in administration of the enterprise. In his subsequent work, Chandler showed how scale and scope economies provided new growth opportunities for the corporation during the second industrial revolution.[2]

Igor Ansoff built upon ideas from *Strategy and Structure* and applied them to emerging concepts of corporate strategy.[3] Strategy came to be seen as a conscious plan to align the firm with opportunities and threats posed by its environment. Kenneth R. Andrews was one of the first theorists to differentiate between a business strategy and a corporate strategy. He held the former to be "the product-market choices made by division or product line management in a diversified company."[4] Corporate strategy was a superset of business strategy. "Like business strategy, [corporate strategy] defines products and markets—and determines the company's course into the almost indefinite future. . . . A company will have only one corporate strategy but may incorporate into its concept of itself several business strategies."[5] Thus, a firm's current businesses influenced its choice of likely future businesses as well, an important insight for understanding corporate innovation.

The subsequent analysis of competitive strategy owes a great deal to the seminal work of Michael Porter. In his first book on the topic,[6] Porter articulated

Henry Chesbrough is the Executive Director of the Center for Open Innovation at the Haas School of Business, University of California, Berkeley. <chesbrou@haas.berkeley.edu>

Melissa M. Appleyard is an Ames Professor in the Management of Innovation and Technology at the School of Business Administration at Portland State University. <MelissaA@sba.pdx.edu>

a conception of strategy that was rooted in the economics of industrial organization, particularly the model of "structure, conduct, and performance."[7] Essentially, Porter cleverly turned Joe S. Bain's economic welfare analysis of monopoly and oligopoly on its head. Instead of maximizing *consumer* surplus (the usual economic objective), Porter focused attention upon those actions that would maximize *producer* surplus. The Porterian model of the Five Forces that shape a firm's competitive strategy—namely, rivalry, buyer power, supplier power, substitutes, and barriers to entry—provided a handy way to identify actions that could enhance a producer's surplus. Items that were previously associated with anti-competitive social welfare outcomes in traditional economic industrial organization theory, such as high barriers to entry, were transformed by Porter's analysis into managerial actions that could enhance a firm's competitive strategy.

In his second volume on strategy, Porter extended the Five Forces concept by linking it to the value chain of a firm, defined as those activities from raw materials through to the final consumer in which a firm's products were developed and marketed.[8] Positions within the value chain in which there were few competitors or other advantageous characteristics (as defined by the above Five Forces model) could create competitive advantage by profiting from other parts of the value chain in which greater competition could be found.

These seminal contributions made an enormous impact upon both the theory and the practice of strategy. With regard to the latter, consulting firms such as McKinsey, Booz Allen, BCG, and Bain soon developed practices and tools that adapted the Porterian notions of strategy for their clients. Porter even launched his own strategy consulting practice, Monitor Company, to apply his strategy concepts for a variety of clients. Monitor continues to enjoy a thriving practice to this day.

Academics also responded to this new approach to strategy in at least four important ways.[9] First, scholars such as Anita McGahan extended Porter's concepts through extensive empirical research that broadly supported Porter's concepts.[10] Second, a former student of Porter's, Richard Rumelt, focused strategy away from industry characteristics toward the characteristics of individual firms. He found that the industry-level differences highlighted in the five forces model were actually less predictive of firm profitability than were differences between firms within a single industry.[11] Third, a related stream of scholarship called the resource-based view of the firm looked within firms to identify the sources of superior firm profitability, and it isolated ownership of certain key resources as the locus of competitive advantage, rather than the Porterian view of a firm's position in its market and its value chain.[12] Finally, a fourth stream examined the role of economic complements to the firm's own assets. Controlling key complementary assets afforded firms a comparative advantage, which facilitated entry into new industries.[13]

Each of these directions has proven to be fruitful for understanding business strategy. None, however, in our judgment, can adequately account for some of the new empirical phenomena emerging in many technology-based industries. All of the traditional views are based upon ownership and control as the key levers in achieving strategic success. All focus largely within the firm, or within the value chain in which the firm is embedded. None take much notice of the potential value of external resources that are not owned by the firm in question, but may nonetheless create value for the firm. These external resources, such as volunteer contributors, innovation communities and ecosystems, and surrounding networks represent growing sources of value creation.

Emerging Anomalies that Challenge Traditional Business Strategy

As Donald Stokes observed, science often progresses first from a practical knowledge of how to do something, to a deeper knowledge of why that something works the way it does.[14] In Thomas Kuhn's notion of paradigm development, empirical anomalies accumulate that (sooner or later) challenge the prevailing conception and trigger the search for an alternative conception that can incorporate the previously inexplicable anomalies.[15] In strategy, we believe that a number of new and anomalous developments have emerged that require a substantive revision to Porter's conceptions, and to the four branches of research that Porter's work has spawned.

While it is difficult to precisely define the scope of these new developments, we believe that the concept of open source development and similarly inspired ideas such as open innovation, the intellectual commons, peer production, and earlier notions of collective invention represent phenomena that require a rethinking of strategy.[16]

Shifting the focus from ownership to the concept of openness requires a reconsideration of the processes that underlie value creation and value capture. Our notion of openness is defined as the pooling of knowledge for innovative purposes where the contributors have access to the inputs of others and cannot exert exclusive rights over the resultant innovation. In its purest form, the value created through an open process would approach that of a public good.[17] It would be "non-rival" in that when someone "consumed" it, it would not degrade the experience of a subsequent user.[18] It also would be "non-excludable" so all comers could gain access.

Typically public goods have been the purview of governments—national defense and education being two widely deployed examples. Recent private-sector phenomena ranging from social networking web sites such as MySpace to open source software such as the Linux operating system have created value along the lines of a public good in that multiple people can use them and no one is excluded from using them.

The value of openness is actually enhanced with every user in two ways. First, users directly contribute ideas and content to improve the quality and variety of the product. MySpace relies on individual contributors, Wikipedia relies on individuals for both data entry and editing, and Linux relies on a global innovation community. Raymond popularized this notion through "Linus's Law," which states, "Given enough eyeballs, all bugs are shallow" (i.e., easy to fix). Second, the more users, the more momentum behind the product such that other companies producing complementary goods or services would be attracted to the mass of users. This dynamic, where more users beget more users, has been labeled a "network effect."[19] In the case of MySpace, Rupert Murdoch's News Corporation found value in the web site's ability to outpace other social networking sites in terms of membership whose demographics—in addition to numbers—are coveted by advertisers.[20] News Corp.'s $580 million acquisition of MySpace's parent company in 2005 put a dollar figure on the value created. The value of Linux's contributions to global computing is reflected in the value of its ecosystem (including software and servers), which was estimated to reach roughly $18 billion in 2006.[21]

These types of open innovation products challenge some of the basic tenets of traditional business strategy. The first tenet called into question is the need to have ownership over the resources that are creating the value. MySpace, YouTube, Wikipedia, and Linux have relied primarily on external, volunteer contributors. The second tenet is the ability to exclude others from copying the product. While ownership of the posted content in the case of MySpace and YouTube certainly is central to their valuations, the users can access the sites and view the content without a charge. Like Linux, Wikipedia relies on its user base to continually refine the product. To guarantee transparency of the open innovation process, Wikipedia has a formalized paper trail whereby the Wikipedia Foundation maintains a log of all of the data entries and the editors of those entries, so that the community can see the origins of entries and the history of subsequent edits to those entries.[22] In the case of Linux, its rules governing the software ensure that the source code will be open for all to see and that the open source code ensures that the kernel will be open for all to see, and that any accepted revisions and improvements will also be open.

When considering the tenets of Porter's Five Forces as the basis of an advantageous competitive position, additional empirical anomalies have emerged. Google and YouTube came into existence without the benefit of significant entry barriers. When considering switching costs on the Web, people can shift to alternative technologies with the click of a mouse. In Porter's view, rivalry reduces industry profits, yet the search industry has many competing technologies with highly profitable companies such as Google and Yahoo! Indeed, Microsoft's masterful cultivation of the Five Forces of Porter has done little to slow Google's meteoric rise in market capitalization. YouTube's acquisition by Google in 2006 for $1.65 billion in stock similarly attests to the fact that entry, even when entry barriers are low, can lead to a formidable value creation.

Towards a More Open Approach to Strategy

Individually, these examples might seem to be mere curiosities. Taken together, though, they imply that something new is going on; something that cannot adequately be explained through the classic conceptions of business strategy. Items that were of central importance in earlier strategy treatments, such as ownership, entry barriers, switching costs, and intra-industry rivalry are of secondary importance in the genesis of the above phenomena. Forces that were either peripheral to the earlier treatment or ignored entirely, such as attracting the participation of individual volunteers, the role of community participation, the construction of innovation networks, and the notion of innovation ecosystems all lay beyond the explanatory power of current notions of strategy.

To further understand value creation and capture in this context, we consider two primary manifestations of openness—open invention and open coordination.

Knowledge Creation through Open Invention

As alluded to above, the power of openness in terms of value creation resets largely with the inherent characteristic of knowledge—it can be reused and can lead to increasing returns.[23] Furthermore, both the breadth and depth of the pooled knowledge can outstrip the knowledge endowment of an individual contributor. One strategic issue for a firm or organization is how to cover the costs of knowledge creation to get this virtual cycle going.

What has proven astounding is that, without direct monetary compensation, a vast number of resources have been committed to open invention, which applies our notion of openness (defined above) to the creation of a new product or service. The poster child for open invention is Linux. Countless person-hours around the globe have been committed to the development, testing, and adoption of this operating system. Skilled programmers rallied around the initial code supplied by Linus Torvalds, and these lead users drove the Linux movement.[24] The enthusiasts that triggered the movement gave rise to an innovation community. The resultant OS has been lauded for its superiority over competing "closed" operating systems along the lines of security, configurability, and reliability.[25] The created value is reflected in the extensive adoption of Linux, where the Linux OS constituted over 13 percent of worldwide server revenue by 2007[26] and has surpassed the Mac OS as the second most widely deployed personal computer OS.[27]

Ecosystem Creation through Open Coordination

In addition to open invention, open coordination has led to consensus building around issues such as technology standards that have permitted whole business ecosystems to flourish. A business ecosystem represents the interplay between multiple industries,[28] so a decision to open up a segment of one industry can have widespread reverberations. As Moore observers, an example from the 1980s is IBM's decision to open up its personal computer (PC) architecture.[29] This led to the rise of the "clones" as companies such as Compaq emulated the IBM specifications. IBM's architecture couple with Microsoft's operating system and Intel's microprocessors became the de facto technology standards in the PC industry.

FIGURE 1. Open and Closed Innovation

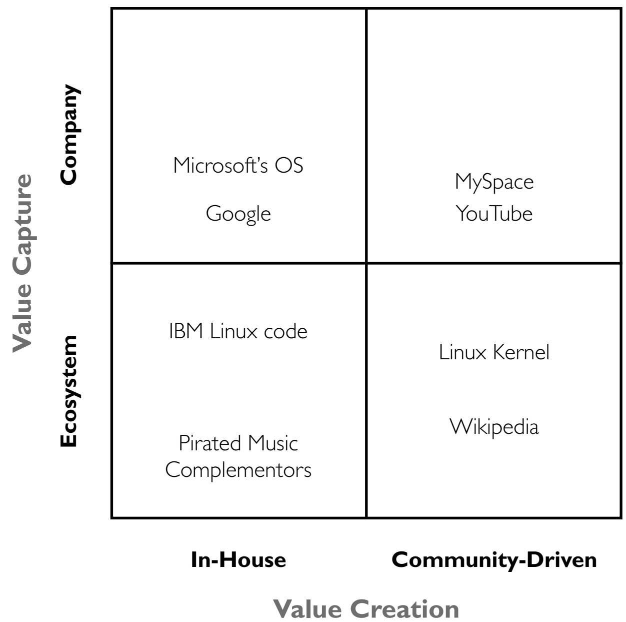

The widespread adoption of this triad contributed to the health of the surrounding ecosystem, which includes application software vendors, video content developers, Internet services providers, and so on. Because PC users want to interact through file sharing and through using numerous software programs, they gravitate to the architecture with the largest footprint. This means that a healthy ecosystem can further perpetuate the adoption of the open architecture through network effects,[30] where the value of the user network is heightened with each additional adopter. Advancing the ecosystem similarly requires community investment in creating new knowledge and exploring alternative architectures to connect the disparate elements of that knowledge together in cohesive ways.[31]

The lingering questions for the business strategist are: Who actually is capturing the value created by open invention and coordination? How are they doing it? The matrix in Figure 1 arrays open initiatives and closed initiatives to illustrate the range of outcomes on both dimensions. On the value creation dimension, initiatives can differ in whether value is created in-house or via a community. On the value capture dimension, an initiative might see its value realized by a company, or by the larger community.

A particular company involved in the innovation process might be able to capture the bulk of the value by closing off the innovation and protecting it with intellectual property (IP) rights—for example, Microsoft's source code for its operating system. Similarly in Google's case, while it captures value from advertisers rather that its user-base, it has been able to distinguish itself through proprietary search algorithms and auction-bidding systems for advertisers. While significant value has accrued to these individual companies, they also have created value that has been captured by their surrounding ecosystems, hence they are placed in the lower portion of the top left quadrant. For example, throu gh its association with Microsoft's operating system, Intel has garnered the leading position in the semiconductor industry, and the personal computer ecosystem has revolved around the "Wintel" de facto standard. By placing paid ads to the right of search results on Google, eBay has bolstered its leadership position in online auctions in the e-commerce ecosystem.

In contrast, in the lower right quadrant, community-driven initiatives can result in products more akin to a public good, leading to value capture that is diffused across an ecosystem. The Linux kernel and Wikipedia are examples. They represent instances of collective invention and coordination. MySpace and YouTube reside in the upper right quadrant, because they rely on community-contributed content, but the IP controls permit the owners of the content, News Corp. and Google, respectively, to "monetize" the content through vehicles such as targeted advertising.[32] The final quadrant, the lower left, reflects innovation initiatives that are fueled by resources within a particular company, but the broader ecosystem captures most of the value, relative to the originator. Two examples populate this quadrant—pirated music and IBM's Linux code. While the proceeds of legitimate music sales accrue to the record labels and their artists and bolster the sales of complementary products in their ecosystem, pirated music only benefits the complementors such as Apple and others, which sell music players. The contribution of code to the Linux kernel by IBM comes from software developers on the payroll of IBM. While IBM can capture value by supplying other goods and services in the value chain, the members of the broader computing ecosystem are free to use the resultant operating system.

A critical element to coordinating the value created through open invention is some underlying architecture that connects the different pieces of knowledge together. This systems-level knowledge may reside in a single company (e.g., IBM in PCs), a collection of firms (e.g., Intel and Microsoft in PCs), a consortium (e.g., SEMATECH in semiconductor equipment), or a nonprofit body (e.g., the Linux Foundation). Without some sense of how the system must operate, open knowledge will not accumulate into useful solutions to real problems.

Open Business Models in Open Source Software

By pooling intellect in a system architecture, open invention and open coordination can produce superior products and services relative to those produced by a smaller number of minds huddled together in a single company. The strategic issue becomes how to capture and then sustain the created value without alienating the individuals, communities, or ecosystem members responsible for the continued development of the good, service, or standard.

While open initiatives often arise from highly motivated individuals or creative communities, a number of approaches have emerged from firms engaged in open innovation to foster value capture and sustainability. Perr, Sullivan, and Appleyard have identified seven "open business models" in the context of open source software (OSS): support, subscription, professional services, proprietary extensions, dual license, device, and community source.[33] In that support, subscription, and professional services are business models found in the proprietary side of the software industry as well, they have not raised many eyebrows. Examples of companies pursuing these models in the open source setting are JBoss (support for application servers),[34] Red Hat (subscriptions for enterprise-versions of Linux), and IBM (a range of professional services for installation and optimization).

Business models novel to the open source software arena include the development of proprietary extensions or add-ons. Companies pursuing this type of model generally have claim to the primary intellectual property covering the application, but they choose an open source software license to help proliferate the product and then offer "enterprise" versions to paying customers, and these versions are generally more stable or have increased functionality. In customer relationship management applications, SugarCRM follows a business model of this sort. The dual license approach is similar to the proprietary extensions model, but it focuses on the type of license under which the software is being distributed. Companies such as MySQL, known for its database products, follows this model by licensing their products under different licenses depending on the intent of the end-user.

The final two business models also are specific to the OSS world. The device model leads companies such as Mazu Networks to offer devices that interact with open source software. In the case of Mazu Networks, the devices are related to network security. The community source model entails having users with almost identical needs pool their resources to address the particular need. The Sakai project pursues collaboration tools for learning environments, and numerous universities are actively involved.

These models can be further grouped in to four categories: deployment, hybridization, complements, and self-service (as reflected in Table 1). In the first category, *deployment* (which spans support, subscription, and professional services), innovation activities heighten the user experience, and users are willing to pay for it even if the initial technology is free.

The second is *hybridization*, in which proprietary innovation investments are made that rely on intellectual property ownership for add-ons (proprietary extensions). A separate instance of this is "versioning,"[35] where multiple versions of a technology such as a public free version and a private commercial version are offered. In open source software, this is called a dual license provision.

The third category is *complements*, where a vendor may sell a PDA, cell phone, or other device at a profit that runs an open source application software suite or operating system. In this category, the value of the complement is actually enhanced by the free nature of the open technology. As the price of the open technology declines, the price to the consumer of the bundled solution

TABLE 1. Open Source Software Business Models

Category	Model	Description	Example
Deployment	Support	Revenue derived from sale of customer support contracts.	JBoss
	Subscription	Revenue derived from annual service agreements bundling open source software, customer support and certified software updates delivered via Internet.	Red Hat Enterprise Linux
	Professional Services/ Consulting	Revenue derived from professional services, training, consulting, or customization of open source software.	IBM
Hybridization	Proprietary Extensions	Firms broadly proliferate open source application and monetize through sale of proprietary versions or product line extensions. Variants include mixed open source/proprietary technologies or services with free trial or "community" versions.	SugarCRM
	Dual License	Vendor licenses software under different licenses (free "Public" or "Community" license vs. paid "Commercial" license) based on customer intent to redistribute.	MySQL
Complements	Device	Vendor sells and supports hardware device or appliance incorporating open source software.	Mazu Networks
Self-Service	Community Source	Consortia of end user organizations or institutions jointly develops application to be used by all.	The Sakai project

Source: Adapted from Jon Perr, Patrick Sullivan, and Melissa M. Appleyard, "Open for Business: Emerging Business Models for Open Source Software Companies," working paper, Lab2Market, Portland State University, 2006.

(open technology plus the complementary device) also falls, thus increasing demand for the device without the manufacturer lowering the price of the device.

The fourth category is a *self-service* model, where a user community creates a software application for its own needs.[36] The first three categories clearly incorporate an element of value capture. Only the last category omits an explicit value capture mechanism. This raises the question of whether this last model is sustainable over time.

These four types of open business models are not mutually exclusive, they may evolve over time, and companies frequently pursue more than one simultaneously. Even firms that have followed the prescriptions of traditional business strategy by placing IP ownership in the center of their business models may wish to consider these approaches to value capture. While a growing number of open invention examples like Linux provide legitimate paths to knowledge creation through volunteerism, an illegal path also exists—piracy. Greatly facilitated by technological change, pirated music and video downloads and knock-off goods (ranging from handbags to pharmaceuticals) have entered the marketplace against the wishes of the original inventors. The enforcement of IP rights can curb the pirates' ability to profit form this "forced" openness, but such legal actions are costly. Because of the difficulty policing and punishing such activity, inventors who thought their business model would rely on patents or copyrights also may wish to consider these alternative approaches to value capture beyond IP enforcement.

Open Innovation beyond IT

The emerging anomalies are by no means confined to the information technology sector. There are a number of new developments in the life sciences, such as the Public Library of Science, where open initiatives are powerfully shaping the face of drug development. This is particularly true for developing new drugs in areas that have not attracted significant commercial interest, such as anti-malarial drugs as well as vaccines. Other recent scholarship has documented the role of innovation communities in the emergence of the snowboard, windsurfing, and skateboarding industries.[37] While we do not wish to suggest that this open approach will migrate to every industry, its emergence is more broad than might be initially realized. As communication costs continue to plummet, facilitating open invention and coordination, it is likely that further open initiatives will take root in more industries around the world.

Issues Confronting the Sustainability of Open Source and Related Initiatives

There are many issues and challenges that the practitioners of increased openness face as they seek to sustain their businesses. While the many successes of open source and related initiatives are rightly acknowledged by their enthusiasts, there are signs that these new approaches to innovation face significant challenges as well. In particular, it is not yet obvious whether and how these initiatives will be able to sustain the ideals and institutions that were used to construct them at the outset. Unless these initiatives demonstrate the ability to prosper and endure, they could become flashes in the pan that, while interesting, ultimately make little impact on technology and society.

Let us start here by examining the single best known and perhaps most successful instance of an open approach: Linux. This open source operating system software was first developed in 1991 by Linus Torvalds. Starting at a code base of roughly 10,000 lines, by 2003, nearly 6 million lines made up the heart of the Linux OS—the Linux kernel. Its support by an extended community is impressive, with more than 130,000 people actively contributing to its development.[38] Linux's market share in network server market is substantial, with a share of 33% in 2007, along with a more modest 3 percent of users in the personal computer segment.[39]

Linux development has been institutionalized through the creation of the Open Source Development Labs (OSDL), located in Portland, Oregon. OSDL was funded largely by the contributions of corporations such as IBM, Intel, HP, and Oracle, who have embraced Linux as part of their own business models. Recently, OSDL merged with the Free Standards Group to form the Linux Foundation, and in our view this merger reflects the success of open source on one hand and its shortcomings on the other.

In terms of success, the merger has been viewed as a testament to the maturity of Linux where consolidation of Linux efforts to assist with issues such as version compatibility was an appropriate next step.[40] Linux has become so successful and so widely adopted that questions of version compatibility have become important. On the side of shortcomings, it was apparent that if OSDL had tried to migrate to a self-funding model by "monetizing" open source opportunities that complemented Linux, its sponsoring corporations might have resisted. This suggests that openness may have a limit if adjacent areas of business are viewed as areas of competition rather than cooperation by corporate sponsors. On the board of the Linux Foundation are again IBM, Intel, HP, and Oracle. Board seats reportedly involve a contribution to the Linux Foundation of $500,000, an amount obviously well beyond an individual's wherewithal that effectively skews the governance of the Linux Foundation towards corporations.[41]

While it is premature to judge the final impact of this restructuring, one can already observe a significant retreat from the initial ideals of the Linux movement, as individuals play a diminished role in the ongoing governance of Linux and corporations play an increasingly important and visible role. One also can infer that a significant risk now exists, where the future development of Linux may be co-opted by the agendas of its corporate governors, rather than the ideals of a community-based meritocracy (in which the best code always wins). One can further infer that the risk is not simply that the Linux agenda may be hijacked; all that is required is that a substantial portion of the community *begins to believe that the agenda is being hijacked*. Once they perceive that to be true, these contributors will take their ideas and contributions elsewhere. This could trigger a collapse within the community, and indeed at that point the corporations would be forced to either support it themselves (thus fulfilling the prophecy) or to abandon it and search for greener pastures.

Thus, the first important issue that open-oriented organizations must face is *how to attract the participation of a broad community of contributors, and then how to sustain their participation over time*. These contributors do not work for the organization and have many other alternative ways to spend their time and talent. If and when a substantial portion of the contributor community perceives that their initiative no longer is driven by the goals that attracted them to the community in the first place, there is a real possibility of collapse of that community.[42]

Linux, we hasten to note, is arguably the most successful example of open source software development. More pedestrian initiatives face considerably more daunting prospects for sustainability. On SourceForge.net, for example, one can find tens of thousands of projects that intend to use an open source method for software development. A casual visit to the site, however, reveals that a few dozen at most have received any significant support from individual software contributors. This reveals a second important issue: the supply of such contributors is not infinite, and the vast majority of projects suffer from a lack of contributors. So *open-oriented projects must compete for contributors*—and most do not succeed in this competition.

One way to compete for contributors is to look for large groups of contributors who can engage with the community. Many such groups can be found inside corporations. In many open source projects, much of the development is done by programmers on the payroll of large corporations.[43] The community contributes to a point and may help with quality control, but company employees contribute the vast majority of the code. This additional participation benefits the open initiative, but raises risks.

A third important issue is *how the open invention or coordination project is led, and how its agenda evolves.* Every community has insiders and outsiders, whether literal or virtual. The insiders typically lead the community and control the direction of its agenda. Most open innovation communities conceive of themselves operating as a meritocracy, where contributors—who often are users of the output as well[44]—provide their inputs for the betterment of the project, as measured by the achievement of the goals and ideals of the project that caused the contributors to join the project initially. If the community becomes dominated by individual contributors who are working for corporations, the perception of a meritocracy rapidly erodes. A sustainable approach to utilizing an innovation community of contributors must identify ways to recruit contributors, keep them engaged, and avoid the perception (let alone the reality) of being co-opted by agendas at odds with the values of that community. In some of the other open examples proffered by enthusiasts such as von Hippel[45] and Shah[46] (such as skateboarding, snowboarding, windsurfing, and the electric guitar industry), innovation started out in open communities but later migrated to become for-profit industries as the number of users grew and a commercial market developed.

A final strategic concern comes from looking at open initiatives from the perspective of the corporation. How can a company engage in an open source community (so as to obtain the benefits of the depth, variety, and quality of technology found in open initiatives) and still profit from that technology, which, by the terms of the intellectual property that governs the community, cannot be owned by the company? *If companies cannot find ways to profit from their innovation activities in open initiatives—through deployment, hybridization, complements, or self-service, they cannot sustain their participation in those initiatives over time.*

While many open source software companies have actively sought community input, over time, the majority of code comes to be written by programmers on staff. This migration from the pure form of open invention to a more hybridized form of open and owned invention is one way that open-oriented firms can control their own destiny. The challenge is managing the mix to avoid alienation of the community, which could precipitate a product war where an open alternative is created to displace the portion that is protected by IP. Well aware of the threat of backlash, open source software companies have been known to focus on developing proprietary code protected by IP only for add-ons that lay outside the areas of interest of the coders in their open innovation community.[47] Clear communication with the open innovation community, confirming that a particular add-on would not be a priority of the community, becomes a managerial imperative.

How Traditional Business Strategy Can Inform Open Initiatives

Ironically, we believe that the best chance for open initiatives to sustain themselves will come from returning to the perspectives of traditional business strategy. If we must compete for contributors to build effective innovation communities, how can we position ourselves to win in that competition? How do we differentiate ourselves to these contributors? If companies must find ways to profit from their participation in open source initiatives, how can they differentiate their products and services in the eyes of customers? Are there places in the value chain or in the surrounding ecosystem where we should be more closed, even as we strive to be open in other places? Are there new business models that combine the prospect of the value creation that derives from openness, with the mechanisms for some degree of value capture necessary for sustainability?

For starters, traditional business strategy has spotlighted settings in which cooperation would likely break down. Fierce rivalry may lead to opportunistic behavior during either open invention or coordination. Alliance partners have been found to engage in "learning races" where the relationship dissolves after one partner aggressively extracts knowledge from the other partner.[48] As dictated by the resource-based view of the firm, employees who are intellectual powerhouses may be jealously guarded, such that their employers would only send "second-stringers" to open invention or coordination initiatives.[49] This could lead to an inferior outcome from the open process. These issues are particularly salient in "one-shot" open initiatives where reputation effects cannot be relied on to deter bad behavior. Mindful of these types of scenarios, leaders of open initiatives can work to establish norms and rules governing the contributors to avoid sub-optimal outcomes for the community.

Traditional strategy also provides two guideposts for value capture. The first points to IP ownership and the second points to creative management of the value chain. As noted above, open source software companies that follow a hybridized business model participate in open invention but also offer either proprietary extensions or a commercial version of their software. At times, this mix between open and closed requires managerial finesse vis-à-vis the community, but in general it has been accepted as a path to profitability. In the case of social networking sites such as MySpace, access may be open, but News Corp.'s ownership of the posted content facilitates additional business opportunities such as a data-mining capability to help with targeted advertising. With additional opt-in features that users are invited to provide for some personal benefit, social networking sites can deliver highly qualified targets for a variety of business purposes.

Even the Porterian notion of the value chain can unleash openness. For example, Intel and IBM have been avid supporters of Linux. Opening up the software link in the electronics value chain has brought down the cost of computing leading to market growth, which means more chip sales for Intel and more hardware sales and service engagements for IBM.[50] Mirroring some OSS companies' sale of devices (as noted above), Intel and IBM sell goods and services that complement the open link in their value chain. Open coordination similarly has "opened up the stack" whereby coordination around interface standards has dismantled monolithic "vertical" value chains like in the telecommunications industry in favor of a bunch of "horizontal" firms specializing in one link of the chain.

Finally, open initiatives may allow for the creation of whole new complementary links in a value chain. As an example, Tim O'Reilly through O'Reilly Media has established a publishing empire in concert with the rise of open source software. The international conferences sponsored by O'Reilly Media are well attended by the OSS faithful, and because he has been so successful in convening intellect, the attendees do not appear to begrudge him his success.

Another strategic perspective that needs to be confronted is whether and when the costs of openness exceed the benefits of openness. Can there be such a thing as too much openness? While more openness is always better in the enthusiasts' accounts of open initiatives, other academic research has found costs, as well as benefits, to developing and maintaining communities and networks. Hansen's analysis of internal networks inside a large firm found that it was costly to maintain ties within the network past a certain size.[51] Laursen and Salter's analysis of data from the British government's Survey of Manufacturers found that respondents' innovation outcomes were positively associated with greater openness (as measured by utilizing a greater number of innovation sources).[52] This association, however, had its limits. Past a certain number of innovation sources, respondents' outcomes became negatively associated with further innovation sources. So more openness and a larger innovation community are valuable, but perhaps only up to a point.

Open Strategy: Illustrative Examples

As we ponder the implications of business strategy for open initiatives, a number of emerging business models attempt to balance the benefits of openness with the need for some value capture for greater sustainability. In addition to the OSS business models noted above, another recent example of an open strategy was the decision of pharmaceutical manufacturer Merck to create the Merck Gene Index. This was an initiative in which Merck funded extensive extramural research activity in universities around the world to produce genetic markers that could serve as targets for later drug development. Once these markers were found, they were compiled and published in Merck's Gene Index. This created a public domain of knowledge that functioned as an intellectual commons for Merck.

While Merck did not have any exclusivity in accessing the markers in its published Index, that was not its objective. Instead, Merck sought to pre-empt the prospect of small biotech firms patenting these markers, thus inhibiting Merck's ability to develop compounds that might turn into new drugs.[53] Merck expected to capture value in its downstream drug development activities and wanted to create a more open source of inputs in the upstream process of identifying potential areas to investigate. So it was balancing value creation upstream in its value chain, while capturing value downstream. This is an instance of what we mean by open strategy.

As noted above, another example of an open strategy that balances value creation and value capture comes from IBM's own involvement with Linux. Readers of a certain age will recall that IBM practiced a distinctly proprietary business model in software for decades, a model that launched products that included Fortran, COBOL, DB2, and AIX, to name but a few of the most salient products. By the late 1990s, however, IBM's software business began to embrace Linux and to construct its own business model around the Linux code. This was a model that was distinctly different from those earlier proprietary software models. As Joel Cawley of IBM explained:

> "I have long observed that it takes $500M to create and sustain a commercially viable OS [operating system]. Today we spend about $100M on Linux development each year. About $50M of that is spent on basic improvements to Linux, how to make it more reliable. The other $50M is spent on things that IBM needs, like special drivers for particular hardware or software to connect with it. We asked the Open Source Development Lab to estimate how much other commercial development spending was being done on Linux. This didn't count any university or individual work, just other companies like us. They told us the number was $800-900M a year, and that the mix of basic vs. specific needs was close to 50/50. So that $500 million investment [required for an operating system] is also there now for Linux as well (counting only the basic portion, not the specific portion). And we only pay $100M toward that. So you can see even from a very narrow accounting view that this is a good business investment for us."[54]

And the specific portion of IBM's funding of Linux allows its internal programmers to optimize the code base to run very effectively with IBM's other hardware and software products. IBM makes good money on these complementary hardware and software items (a variation on the device category noted above), so participating in a community at one level of value creation leads to greater value capture higher up the stack of value added activities for IBM.

Executing this new, open strategy required some major internal changes within IBM, and also required IBM to change the opinions of many outsiders who were skeptical about working with IBM. It wasn't easy. Outside Linux participants, for example, were afraid that IBM would destroy the values of the Linux community, either intentionally or unintentionally. As Jerry Stallings, IBM's VP of IP and Strategy described it, "IBM's reputation was a big sometimes arrogant company that takes over whatever it gets involved in. We had to learn how to collaborate."

Conclusion: Open Strategy Balances Value Creation with Value Capture

Open strategy balances the powerful value creation forces that can be found in creative individuals, innovation communities, and collaborative initiatives with the need to capture value in order to sustain continued participation and support of those initiatives. Traditional concepts of business strategy either underestimate the value of open invention and open coordination, or they ignore them outright. As the concept of openness spreads from software to science and other industries, we will need to update our concepts of strategy. Open strategy is an attempt to supply this update.

In open-dominated industry segments, such as open source software, new business models have been established. The models often blend elements of open and closed innovation. The OSS business models fall under four primary categories: deployment, hybridization, complements, and self-service. These models may apply to other industries as openness spreads.

At the same time, open initiatives must confront real and serious challenges to their ability to sustain themselves over time. While building broad communities of motivated individuals can unleash creative contributions, these are difficult to sustain over time. Attracting and retaining contributors, preventing co-option of the innovation agenda, and covering the fixed costs of innovation all represent non-trivial managerial headaches. As noted, even the most celebrated example of openness, the Linux kernel, now confronts significant changes that may threaten its ability to remain open.

These issues of sustainability bring us back to traditional business strategy, which can make important contributions to mitigating them. If we are to make strategic sense of innovation communities, ecosystems, networks, and their implications for competitive advantage, we propose that a new approach to strategy—open strategy—is needed. Open strategy balances the tenets of traditional business strategy with the promise of open innovation. Certain companies appear to be constructing open strategies. These examples are worth studying, and may point the way forward for both openness and for strategy in leading through innovation.

Notes

1. Alfred Chandler, *Strategy and Structure* (Cambridge, MA: MIT Press, 1962).
2. Alfred Chandler, "The Enduring Logic of Industrial Success," *Harvard Business Review*, 68/2 (March/April 1990): 130-140.
3. Igor Ansoff, *Corporate Strategy* (New York, NY: McGraw-Hill. 1965)
4. Kenneth R. Andrews, *The Concept of Corporate Strategy* (Homewood, IL: Irwin, 1987).
5. Ibid., p. xi.
6. Michael Porter, *Competitive Strategy* (New York, NY: Free Press, 1980).
7. Joe S. Bain, 1956. *Barriers to New Competition* (Cambridge, MA: Harvard University Press, 1956).
8. Michael Porter, *Competitive Advantage* (New York, NY: Free Press, 1985).
9. It is perhaps inevitable that our account of strategy is selective, and omits certain areas of academic inquiry. Among the branches that we have excluded from this account are approaches such as game theoretic views of strategy (Shapiro), sociologically informed approaches to strategy (Baum and Dutton), and chaos-based views of strategy (Brown and Eisenhart). Carl Shapiro, "The Theory of Business Strategy," *The RAND Journal of Economics*, 20/1 (Spring 1989): 125-137; Joel A.C. Baum and J.E. Dutton, eds., *Advances in Strategic Management: The Embeddedness of Strategy* (Greenwich, CT: JAI Press, 1996); Shona L. Brown and Kathleen M. Eisenhardt, *Competing on the Edge : Strategy as Structured Chaos* (Cambridge, MA: Harvard University Press, 1998.
10. Anita M. McGahan and Michael E. Porter, "How Much Does Industry Matter, Really?" *Strategic Management Journal*, 18/6 (Summer Special Issue 1997): 15-30; Anita M. McGahan and Michael E. Porter, "What Do We Know About Variance in Accounting Profitability?" *Management Science*, 48/7 (July 2002): 834-851.
11. Richard Rumelt, "Diversification Strategy and Profitability," *Strategic Management Journal*, 3/4 (October-December 1982): 359-369; Richard P. Rumelt, "How Much Does Industry Matter?" *Strategic Management Journal*, 12/3 (March 1991): 167-185.
12. Birger Wernerfelt, "A Resource-Based View of the Firm," *Strategic Management Journal*, 5/2 (April/June 1984): 171-180; Jay B. Barney, "Strategic Factor Markets: Expectations, Luck, and Business Strategy," *Management Science*, 32/10 (October 1986): 1231-1241; Jay B. Barney, "Firm Resources and Sustained Competitive Advantage," *Journal of Management*, 17/1 (March 1991): 99-120; Ingemar Dierickx and Karel Cool, "Asset Stock Accumulation and the Sustainability of Competitive Advantage," *Management Science*, 35/12 (December 1989): 1504-1511; Ingemar Dierickx and Karel Cool, "Asset Stock Accumulation and the Sustainability of Competitive Advantage: Reply," *Management Science*, 35/12 (December 1989): 1514-1514; Joseph Mahoney and J. Rajendran Pandian, "The Resource-Based View within the Conversation of Strategic Management," *Strategic Management Journal*, 13/5 (June 1992): 363-380; M.A. Peteraf, "The Cornerstones of Competitive Advantage: A Resource-Based View," *Strategic Management Journal*, 14/3 (March 1993): 179-191.
13. David Teece, "Profiting from Technological Innovation," *Research Policy*, 15/6 (1986): 285-305.
14. Donald Stokes, *Pasteur's Quadrant: Basic Science and Technological Innovation* (Washington, D.C.: Brookings Institution Press, 1997).
15. Thomas Kuhn, *The Structure of Scientific Revolutions* (Chicago, IL: University of Chicago Press, 1962).
16. For open source development, see Eric Raymond, *The Cathedral and the Bazaar* (Beijing: O'Reilly, 1999); Bruce Perens, *The Emerging Economic Paradigm of Open Source* <http://perens.com/Articles/Economic.html>, February 16, 2005; Eric von Hippel, *Democratizing Innovation* (Cambridge, MA: MIT Press, 2005). For open innovation, see Henry Chesbrough, *Open Innovation: the New Imperative for Creating and Profiting from Technology* (Boston, MA: Harvard Business School Press, 2003). For the intellectual commons, see Lawrence Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World* (New York, NY: Random House, 2001). For peer production, see Josh Lerner and Jean Tirole, "Some Simple Economics of Open Source," *Journal of Industrial Economics*, 50/2 (June 2002): 197-234; Yochai Benkler, "Coase's Penguin, or, Linux and the Nature of the Firm," *Yale Law Journal*, 112/3 (December 2002): 369. For collective invention, see R.C. Allen, "Collective Invention," *Journal of Economic Behavior and Organization*, 4/1 (March 1983): 1-24.
17. Benkler, op. cit.

18. Paul M. Romer, "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 94/5 (October 1986): 1002-1037.
19. Brian Arthur, *Increasing Returns and Path Dependence in the Economy* (Ann Arbor, MI: The University of Michigan Press, 1994).
20. Patricia Sellers, "MySpace Cowboys," <http://money.cnn.com/magazines/fortune/fortune_archive/2006/09/04/8384727/index.htm>, August 29 2006.
21. Sean Michael Kerner, "IDC: Linux Ecosystem Worth $40 Billion by 2010," <http://www.internetnews.com/dev-news/article.php/3659961>, February 14, 2007.
22. Don Tapscott and Anthony D. William, *Wikinomics: How Mass Collaboration Changes Everything* (New York, NY: Portfolio, 2006).
23. W. Brian Arthur, "Increasing Returns and the New World of Business," *Harvard Business Review*, 74/4 (July/August 1996): 100-110; Paul M. Romer, "Endogenous Technological Change," *Journal of Political Economy*, 98/5 (October 1990): S71-S102.
24. von Hippel, op. cit.; Sonali Shah, "Motivation, Governance, and the Viability of Hybrid Forms in Open Source Software Development," *Management Science*, 52/7 (July 2006): 1000-1014.
25. Michelle Delio, "Linux: Fewer Bugs Than Rivals," <www.wired.com/software/coolapps/news/2004/12/66022>, December 14, 2004; Mark Brunelli, "Users Tackle Question of Linux vs. Windows on the Server," <http://searchenterpriselinux.techtarget.com/originalContent/0,289142,sid39_gci1123175,00.html?topic=300780>, September 8, 2005.
26. IDC, "Server Market Accelerates as New Workloads and a Strong Refresh Cycle Drive Server Demand in the Enterprise," <www.idc.com/getdoc.jsp;jsessionid=CHZBXSXEQGMI2CQJAFDCFFAKBEAVAIWD?containerId=prUS20840807>, August 23, 2007.
27. Robert McMillan, "Analysis: The Business Case for Desktop Linux," <www.computerworld.com.au/index.php/id;1048043952;fp;16;fpid;0>, December 24, 2004.
28. James Moore, *The Death of Competition* (Boston, MA: Harvard Business School Press, 1993)
29. Ibid.
30. Carl Shapiro and Hal Varian, *Information Rules: A Strategic Guide to the Network Economy* (Boston, MA: Harvard Business School Press, 1998).
31. Annabelle Gawer and Michael Cusumano, *Platform Leadership* (Boston, MA: Harvard Business School Press, 2002); Marco Iansiti and Roy Levien, *The Keystone Advantage* (Boston, MA: Harvard Business School Press, 2004).
32. Ownership also makes these sites potential targets for legal actions by copyright owners who feel that the site is monetizing value without paying appropriate compensation for the copyrighted content hosted on the site, such as the recent Viacom suit against YouTube. That is a subject for another paper, however.
33. Jon Perr, Patrick Sullivan, and Melissa M. Appleyard, "Open for Business: Emerging Business Models for Open Source Software Companies," working paper, Lab2Market, Portland State University, 2006; Henry Chesbrough, *Open Business Models: How to Thrive in the New Innovation Landscape* (Boston, MA: Harvard Business School Press, 2006).
34. JBoss was acquired by Red Hat in 2006.
35. Shapiro and Varian, op. cit.
36. von Hippel, op. cit.
37. Shah, op. cit.
38. <http://counter.li.org/>.
39. <http://blogs.zdnet.com/ITFacts/?p=6932>.
40. Mike Rogoway, "Merger Marks Open-Source Milestone," *The Oregonian*, January 23, 2007, pp. C1-C2.
41. Neil McAllister, "Questioning the Linux Foundation's credentials: How noble are its motivations?" CNET News, February 12, 2007.
42. Jared Diamond, *Collapse: How Societies Choose to Fail or Succeed* (New York, NY: Viking, 2005).
43. Linus Dahlander and Martin W. Wallin, "A Man on the Inside: Unlocking Communities as Complementary Assets," *Research Policy*, 35/8 (October 2006): 1243-1259.
44. von Hippel, op. cit.
45. Ibid.
46. Shah, op. cit.
47. Perr et al., op. cit.

48. Gary Hamel, "Competition for Competence and Inter-Partner Learning within International Strategic Alliances," *Strategic Management Journal*, 12/4 (Summer 1991): 83-103.
49. Peter Grindley, David C. Mowery, and Brian Silverman, "SEMATECH and Collaborative Research: Lessons in the Design of High-Technology Consortia," *Journal of Policy Analysis and Management*, 13/4 (Autumn 1994): 723-758.
50. Chesbrough (2006), op. cit.
51. Morten Hansen, "The Search-Transfer Problem: The Role of Weak Ties in Sharing Knowledge across Organization Subunits," *Administrative Science Quarterly*, 44/1 (March 1999): 82.
52. Keld Laursen and Ammon Salter, "Open for Innovation: The Role of Openness in Explaining Innovation Performance among U.K. Manufacturing Firms," *Strategic Management Journal*, 27/2 (February 2006): 131-150.
53. Robert P. Merges, "A New Dynamism in the Public Domain," *University of Chicago Law Review*, 71 (2004): 183-203; Gary Pisano, *Science Business: The Promise, the Reality, and the Future of Biotech* (Boston, MA: Harvard Business School Press, 2006).
54. Cawley's quote is taken from Chesbrough (2006), op. cit., pp. 193-194.

Procter & Gamble's radical strategy of open innovation now produces more than 35% of the company's innovations and billions of dollars in revenue.

CONNECT AND DEVELOP

INSIDE PROCTER & GAMBLE'S NEW MODEL FOR INNOVATION

by Larry Huston and Nabil Sakkab

Procter & Gamble launched a new line of Pringles potato crisps in 2004 with pictures and words – trivia questions, animal facts, jokes – printed on each crisp. They were an immediate hit. In the old days, it might have taken us two years to bring this product to market, and we would have shouldered all of the investment and risk internally. But by applying a fundamentally new approach to innovation, we were able to accelerate Pringles Prints from concept to launch in less than a year and at a fraction of what it would have otherwise cost. Here's how we did it.

Back in 2002, as we were brainstorming about ways to make snacks more novel and fun, someone suggested that we print pop culture images on

NICK LOWNDES

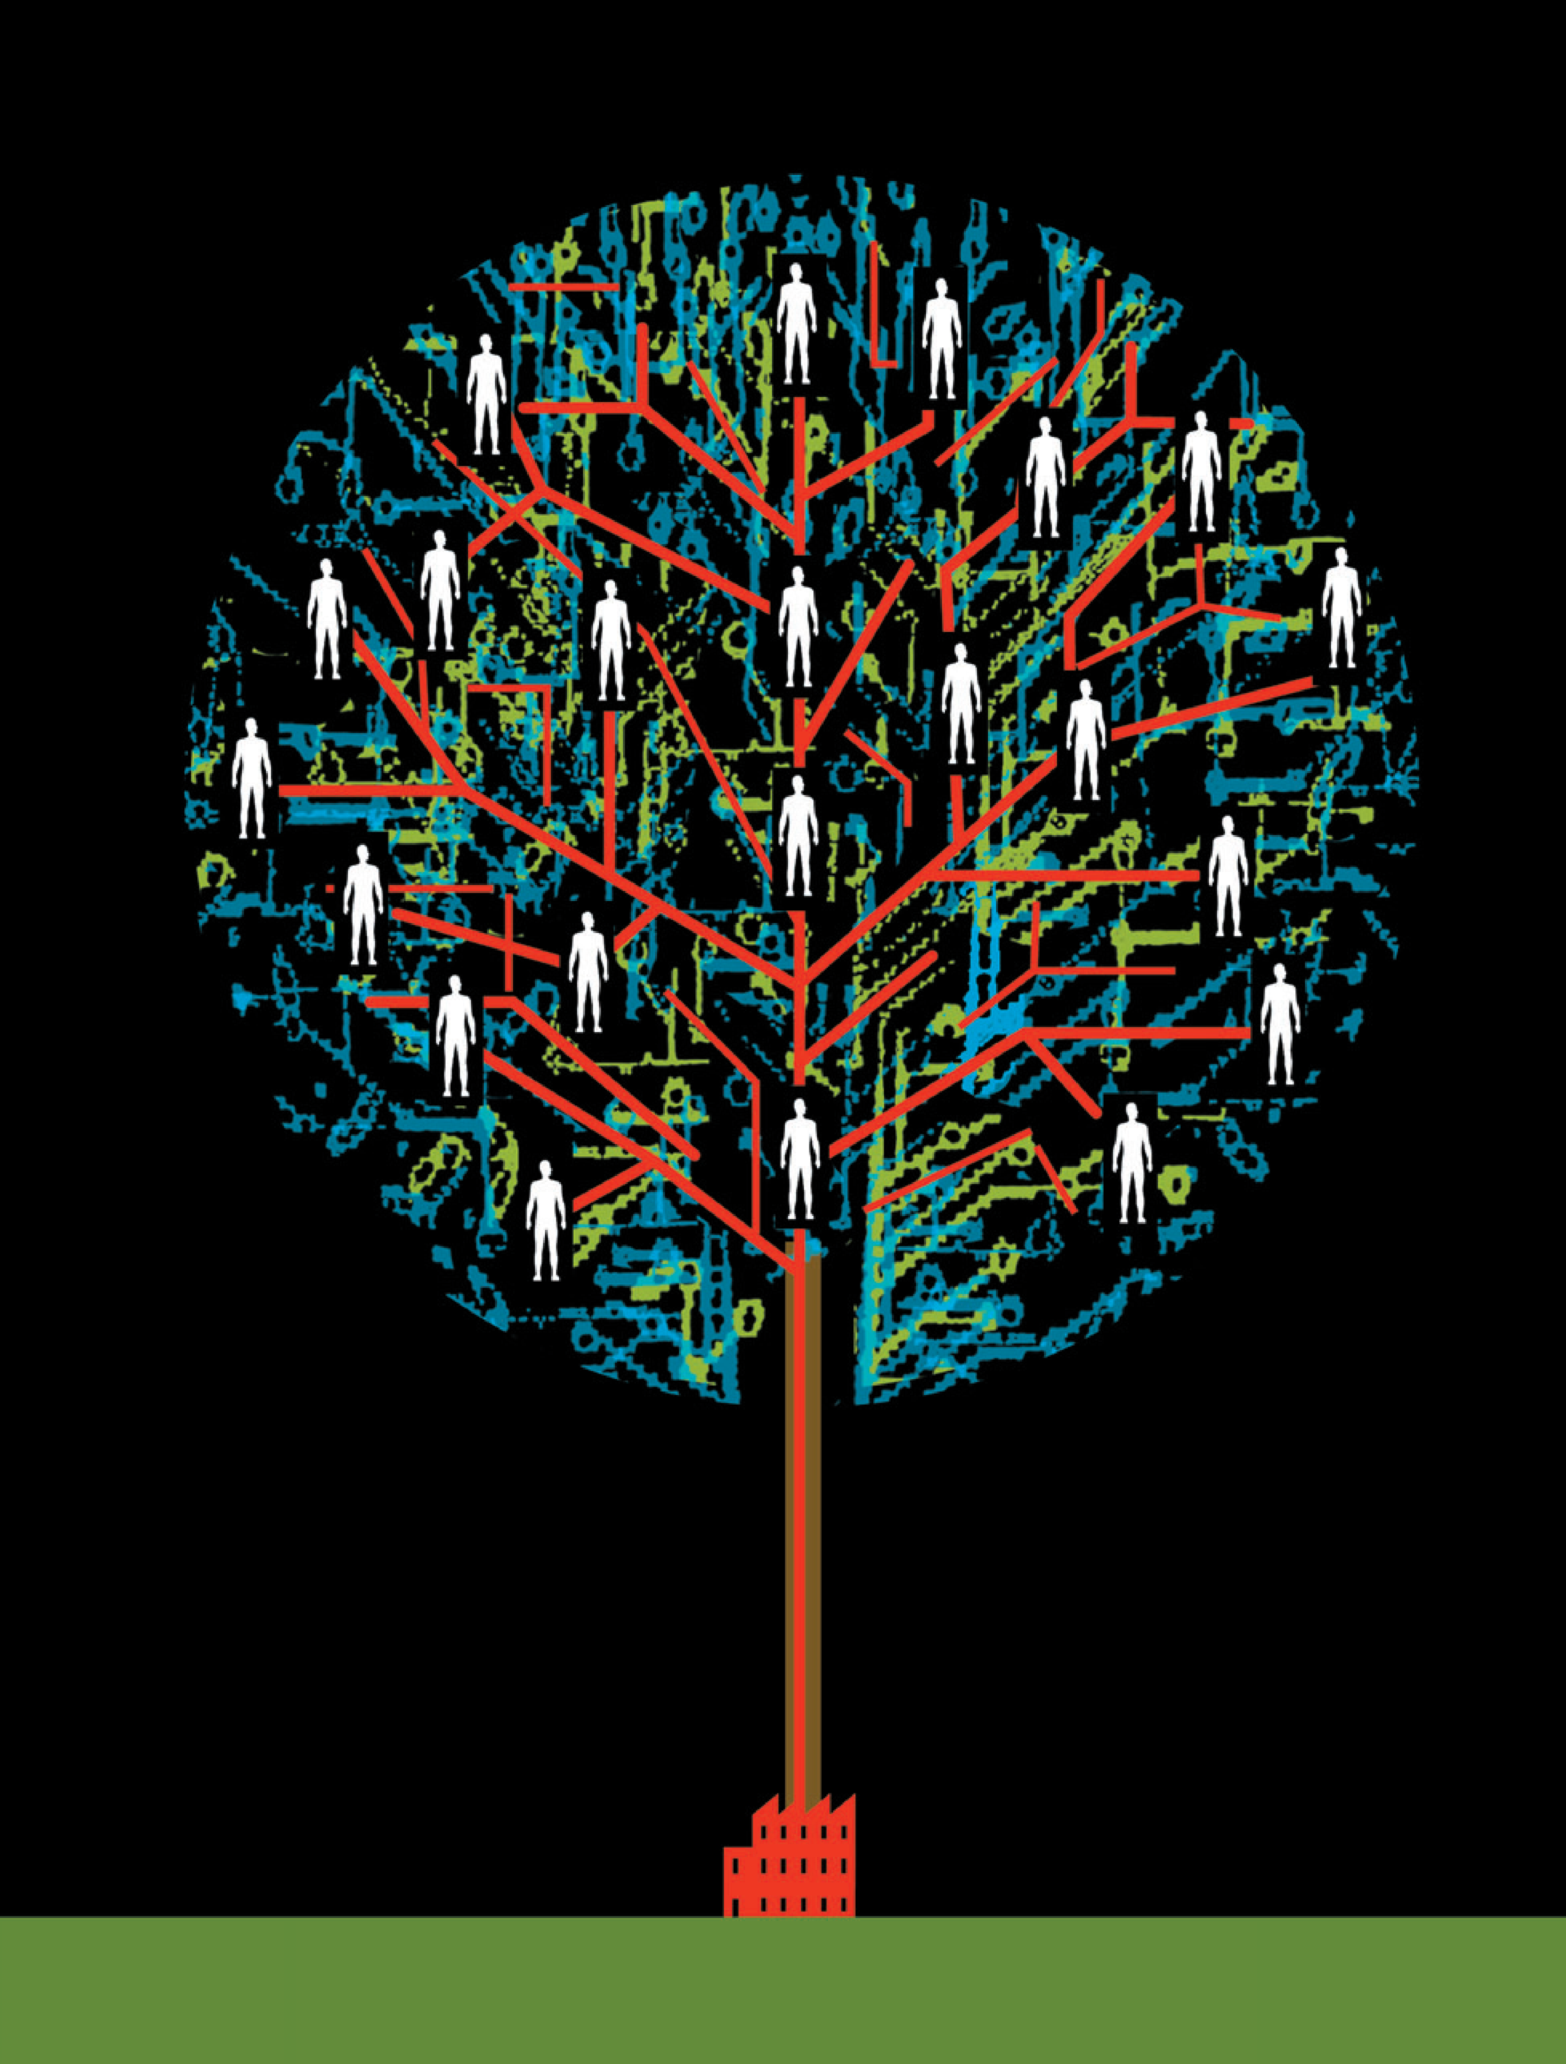

Pringles. It was a great idea, but how would we do it? One of our researchers thought we should try ink-jetting pictures onto the potato dough, and she used the printer in her office for a test run. (You can imagine her call to our computer help desk.) We quickly realized that every crisp would have to be printed as it came out of frying, when it was still at a high humidity and temperature. And somehow, we'd have to produce sharp images, in multiple colors, even as we printed thousands upon thousands of crisps each minute. Moreover, creating edible dyes that could meet these needs would require tremendous development.

Traditionally, we would have spent the bulk of our investment just on developing a workable process. An internal team would have hooked up with an ink-jet printer company that could devise the process, and then we would have entered into complex negotiations over the rights to use it.

Instead, we created a technology brief that defined the problems we needed to solve, and we circulated it throughout our global networks of individuals and institutions to discover if anyone in the world had a ready-made solution. It was through our European network that we discovered a small bakery in Bologna, Italy, run by a university professor who also manufactured baking equipment. He had invented an ink-jet method for printing edible images on cakes and cookies that we rapidly adapted to solve our problem. This innovation has helped the North America Pringles business achieve double-digit growth over the past two years.

From R&D to C&D

Most companies are still clinging to what we call the invention model, centered on a bricks-and-mortar R&D infrastructure and the idea that their innovation must principally reside within their own four walls. To be sure, these companies are increasingly trying to buttress their laboring R&D departments with acquisitions, alliances, licensing, and selective innovation outsourcing. And they're launching Skunk Works, improving collaboration between marketing and R&D, tightening go-to-market criteria, and strengthening product portfolio management.

But these are incremental changes, bandages on a broken model. Strong words, perhaps, but consider the facts: Most mature companies have to create organic growth of 4% to 6% year in, year out. How are they going to do it? For P&G, that's the equivalent of building a $4 billion business this year alone. Not long ago, when companies were smaller and the world was less competitive, firms could rely on internal R&D to drive that kind of growth. For generations, in fact, P&G created most of its phenomenal growth by innovating from within – building global research facilities and hiring and holding on to the best talent in the world. That worked well when we were a $25 billion company; today, we're an almost $70 billion company.

By 2000, it was clear to us that our invent-it-ourselves model was not capable of sustaining high levels of top-line growth. The explosion of new technologies was putting ever more pressure on our innovation budgets. Our R&D productivity had leveled off, and our innovation success rate – the percentage of new products that met financial objectives – had stagnated at about 35%. Squeezed by nimble competitors, flattening sales, lackluster new launches, and a quarterly earnings miss, we lost more than half our market cap when our stock slid from $118 to $52 a share. Talk about a wake-up call.

The world's innovation landscape had changed, yet we hadn't changed our own innovation model since the late 1980s, when we moved from a centralized approach to a globally networked internal model – what Christopher Bartlett and Sumantra Ghoshal call the transnational model in *Managing Across Borders.*

We discovered that important innovation was increasingly being done at small and midsize entrepreneurial companies. Even individuals were eager to license and sell their intellectual property. University and government labs had become more interested in forming industry partnerships, and they were hungry for ways to monetize their research. The Internet had opened up access to talent markets throughout the world. And a few forward-looking companies like IBM and Eli Lilly were beginning to experiment with the new concept of open innovation, leveraging one another's (even competitors') innovation assets – products, intellectual property, and people.

As was the case for P&G in 2000, R&D productivity at most mature, innovation-based companies today is flat while innovation costs are climbing faster than top-line growth. (Not many CEOs are going to their CTOs and saying, "Here, have some more money for innovation.") Meanwhile, these companies are facing a growth mandate that their existing innovation models can't possibly support. In 2000, realizing that P&G couldn't meet its growth objectives by spending more and more on R&D for less and less payoff, our newly appointed CEO, A.G. Lafley, challenged us to reinvent the company's innovation business model.

We knew that most of P&G's best innovations had come from connecting ideas across internal businesses. And after studying the performance of a small number of products we'd acquired beyond our own labs, we knew that external connections could produce highly profitable innovations, too. Betting that these connections were the

Larry Huston *(huston.la@pg.com) is the vice president for innovation and knowledge and* ***Nabil Sakkab*** *(sakkab.ny@pg.com) is the senior vice president for corporate research and development at Procter & Gamble in Cincinnati.*

key to future growth, Lafley made it our goal to acquire 50% of our innovations outside the company. The strategy wasn't to replace the capabilities of our 7,500 researchers and support staff, but to better leverage them. Half of our new products, Lafley said, would come *from* our own labs, and half would come *through* them.

It was, and still is, a radical idea. As we studied outside sources of innovation, we estimated that for every P&G researcher there were 200 scientists or engineers elsewhere in the world who were just as good – a total of perhaps 1.5 million people whose talents we could potentially use. But tapping into the creative thinking of inventors and others on the outside would require massive operational changes. We needed to move the company's attitude from resistance to innovations "not invented here" to enthusiasm for those "proudly found elsewhere." And we needed to change how we defined, and perceived, our R&D organization – from 7,500 people inside to 7,500 *plus* 1.5 million outside, with a permeable boundary between them.

It was against this backdrop that we created our *connect and develop* innovation model. With a clear sense of consumers' needs, we could identify promising ideas throughout the world and apply our own R&D, manufacturing, marketing, and purchasing capabilities to them to create better and cheaper products, faster.

The model works. Today, more than 35% of our new products in market have elements that originated from outside P&G, up from about 15% in 2000. And 45% of the initiatives in our product development portfolio have key elements that were discovered externally. Through connect and develop – along with improvements in other aspects of innovation related to product cost, design, and marketing – our R&D productivity has increased by nearly 60%. Our innovation success rate has more than doubled, while the cost of innovation has fallen. R&D investment as a percentage of sales is down from 4.8% in 2000 to 3.4% today. And, in the last two years, we've launched more than 100 new products for which some aspect of execution came from outside the company. Five years after the company's stock collapse in 2000, we have doubled our share price and have a portfolio of 22 billion-dollar brands.

According to a recent Conference Board survey of CEOs and board chairs, executives' number one concern is "sustained and steady top-line growth." CEOs understand the importance of innovation to growth, yet how many have overhauled their basic approach to innovation? Until companies realize that the innovation landscape has changed and acknowledge that their current model is unsustainable, most will find that the top-line growth they require will elude them.

Where to Play

When people first hear about connect and develop, they often think it's the same as outsourcing innovation – contracting with outsiders to develop innovations for P&G. But it's not. Outsourcing strategies typically just transfer work to lower-cost providers. Connect and develop, by contrast, is about finding good ideas and

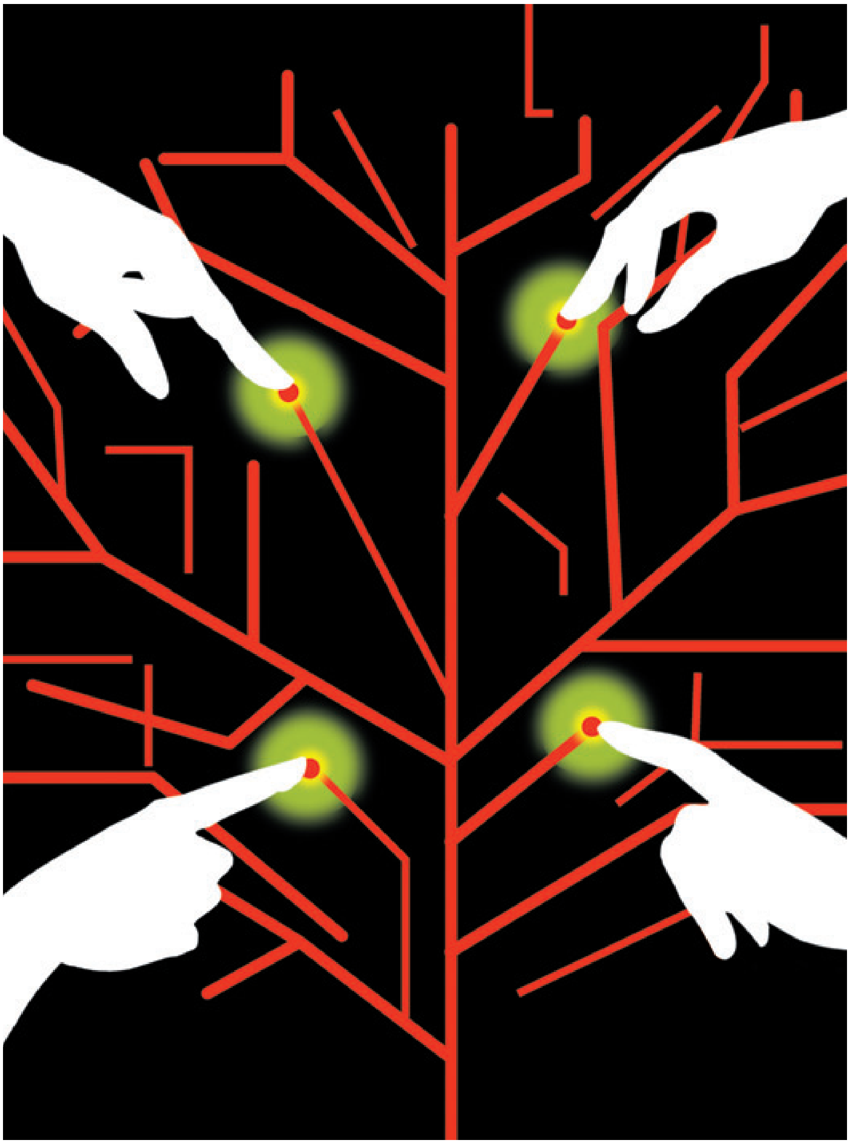

Most companies are still clinging to a bricks-and-mortar R&D infrastructure and the idea that their innovation must principally reside within their own four walls.

bringing them in to enhance and capitalize on internal capabilities.

To do this, we collaborate with organizations and individuals around the world, systematically searching for proven technologies, packages, and products that we can improve, scale up, and market, either on our own or in partnership with other companies. Among the most successful products we've brought to market through connect and develop are Olay Regenerist, Swiffer Dusters, and the Crest SpinBrush.

For connect and develop to work, we realized, it was crucial to know exactly what we were looking for, or where to play. If we'd set out without carefully defined targets, we'd have found loads of ideas but perhaps none that were useful to us. So we established from the start that we would seek ideas that had some degree of success already; we needed to see, at least, working products, prototypes, or technologies, and (for products) evidence of consumer interest. And we would focus on ideas and products that would benefit specifically from the application of P&G technology, marketing, distribution, or other capabilities.

Then we determined the areas in which we would look for these proven ideas. P&G is perhaps best known for its personal hygiene and household-cleaning products – brands like Crest, Charmin, Pampers, Tide, and Downy. Yet we produce more than 300 brands that span, in addition to hygiene and cleaning, snacks and beverages, pet nutrition, prescription drugs, fragrances, cosmetics, and many other categories. And we spend almost $2 billion a year on R&D across 150 science areas, including materials, biotechnology, imaging, nutrition, veterinary medicine, and even robotics.

To focus our idea search, we directed our surveillance to three environments:

Top ten consumer needs. Once a year, we ask our businesses what consumer needs, when addressed, will drive the growth of their brands. This may seem like an obvious question, but in most companies, researchers are working on the problems that they find interesting rather than those that might contribute to brand growth. This inquiry produces a top-ten-needs list for each business and one for the company overall. The company list, for example, includes needs such as "reduce wrinkles, improve skin texture and tone," "improve soil repellency and restoration of hard surfaces," "create softer paper products with lower lint and higher wet strength," and "prevent or minimize the severity and duration of cold symptoms."

These needs lists are then developed into science problems to be solved. The problems are often spelled out in technology briefs, like the one we sent out to find an ink-jet process for Pringles Prints. To take another example, a major laundry need is for products that clean effectively using cold water. So, in our search for relevant innovations, we're looking for chemistry and biotechnology solutions that allow products to work well at low temperatures. Maybe the answer to our cold-water-cleaning problem is in a lab that's studying enzymatic reactions in microbes that thrive under polar ice caps, and we need only to find the lab.

Adjacencies. We also identify adjacencies – that is, new products or concepts that can help us take advantage of existing brand equity. We might, for instance, ask which baby care items – such as wipes and changing pads – are adjacent to our Pampers disposable diapers, and then seek out innovative emerging products or relevant technologies in those categories. By targeting adjacencies in oral care, we've expanded the Crest brand beyond toothpaste to include whitening strips, power toothbrushes, and flosses.

Technology game boards. Finally, in some areas, we use what we call technology game boards to evaluate how technology acquisition moves in one area might affect products in other categories. Conceptually, working with these planning tools is like playing a multilevel game of chess. They help us explore questions such as "Which of our key technologies do we want to strengthen?" "Which technologies do we want to acquire to help us better compete with rivals?" and "Of those that we already own, which do we want to license, sell, or codevelop further?" The answers provide an array of broad targets for our innovation searches and, as important, tell us where we shouldn't be looking.

How to Network

Our global networks are the platform for the activities that, together, constitute the connect-and-develop strategy. But networks themselves don't provide competitive advantage any more than the phone system does. It's how you build and use them that matters.

Within the boundaries defined by our needs lists, adjacency maps, and technology game boards, no source of ideas is off-limits. We tap closed proprietary networks and open networks of individuals and organizations available to any company. Using these networks, we look for ideas in government and private labs, as well as academic and other research institutions; we tap suppliers, retailers, competitors, development and trade partners, VC firms, and individual entrepreneurs.

Here are several core networks that we use to seek out new ideas. This is not an exhaustive list; rather, it is a snapshot of the networking capabilities that we've found most useful.

Proprietary networks. We rely on several proprietary networks developed specifically to facilitate connect-and-develop activities. Here are two of the largest ones.

Technology entrepreneurs. Much of the operation and momentum of connect and develop depends on our network of 70 technology entrepreneurs based around the world. These senior P&G people lead the development of our needs lists, create adjacency maps and technology game boards, and write the technology briefs that define the problems we are trying to solve. They create external connections by, for example, meeting with university and industry researchers and forming supplier networks, and they actively promote these connections to decision makers in P&G's business units.

The technology entrepreneurs combine aggressive mining of the scientific literature, patent databases, and other data sources with physical prospecting for ideas – say, surveying the shelves of a store in Rome or combing product and technology fairs. Although it's effective and necessary to scout for ideas electronically, it's not sufficient. It was a technology entrepreneur who, exploring a local market in Japan, discovered what ultimately became the Mr. Clean Magic Eraser. We surely wouldn't have found it otherwise. (See the exhibit "The Osaka Connection.")

The technology entrepreneurs work out of six connect-and-develop hubs, in China, India, Japan, Western Europe, Latin America, and the United States. Each hub focuses on finding products and technologies that, in a sense, are specialties of its region: The China hub, for example, looks in particular for new high-quality materials and cost innovations (products that exploit China's unique ability to make things at low cost). The India hub seeks out local talent in the sciences to solve problems – in our manufacturing processes, for instance – using tools like computer modeling.

Thus far, our technology entrepreneurs have identified more than 10,000 products, product ideas, and promising technologies. Each of these discoveries has undergone a formal evaluation, as we'll describe further on.

Suppliers. Our top 15 suppliers have an estimated combined R&D staff of 50,000. As we built connect and develop, it didn't take us long to realize that they represented a huge potential source of innovation. So we created a secure IT platform that would allow us to share technology briefs with our suppliers. If we're trying to find ways to make detergent perfume last longer after clothes come out of the dryer, for instance, one of our chemical suppliers may well have the solution. (Suppliers can't see others' responses, of course.) Since creating our supplier network, we've seen a 30% increase in innovation projects jointly staffed with P&G's and suppliers' researchers. In some cases, suppliers' researchers come to work in our labs, and in others, we work in theirs – an example of what we call "cocreation," a type of collaboration that goes well beyond typical joint development.

We also hold top-to-top meetings with suppliers so our senior leaders can interact with theirs. These meetings, along with our shared-staff arrangements, improve relationships, increase the flow of ideas, and strengthen each company's understanding of the other's capabilities – all of which helps us innovate.

Open networks. A complement to our proprietary networks are open networks. The following four are particularly fruitful connect-and-develop resources.

Leading Connect and Develop

The connect-and-develop strategy requires that a senior executive have day-to-day accountability for its vision, operations, and performance. At P&G, the vice president for innovation and knowledge has this responsibility. Connect-and-develop leaders from each of the business units at P&G have dotted-line reporting relationships with the VP. The managers for our virtual R&D networks (such as NineSigma and our supplier network), the technology entrepreneur and hub network, our connect-and-develop legal resources, and our training resources report directly.

The VP oversees the development of networks and new programs, manages a corporate budget, and monitors the productivity of networks and activities. This includes tracking the performance of talent markets like NineSigma and InnoCentive as well as measuring connect-and-develop productivity by region – evaluating, for example, the costs and output (as measured by products in market) of foreign hubs. Productivity measurements for the entire program are reported annually.

NineSigma. P&G helped create NineSigma, one of several firms connecting companies that have science and technology problems with companies, universities, government and private labs, and consultants that can develop solutions. Say you have a technical problem you want to crack–for P&G, as you'll recall, one such problem is cold-temperature washing. NineSigma creates a technology brief that describes the problem, and sends this to its network of thousands of possible solution providers worldwide. Any solver can submit a nonconfidential proposal back to NineSigma, which is transmitted to the contracting company. If the company likes the proposal, NineSigma connects the company and solver, and the project proceeds from there. We've distributed technology briefs to more than 700,000 people through NineSigma and have as a result completed over 100 projects, with 45% of them leading to agreements for further collaboration.

InnoCentive. Founded by Eli Lilly, InnoCentive is similar to NineSigma – but rather than connect companies with contract partners to solve broad problems across many disciplines, InnoCentive brokers solutions to more narrowly defined scientific problems. For example, we might have an industrial chemical reaction that takes five steps to accomplish and want to know if it can be done in three. We'll put the question to InnoCentive's 75,000 contract scientists and see what we get back. We've had problems solved by a graduate student in Spain, a chemist in India, a freelance chemistry consultant in the United States, and an agricultural chemist in Italy. About a third of the problems we've posted through InnoCentive have been solved.

YourEncore. In 2003, we laid the groundwork for a business called YourEncore. Now operated independently, it connects about 800 high-performing retired scientists and engineers from 150 companies with client businesses. By using YourEncore, companies can bring people with deep experience and new ways of thinking from other organizations and industries into their own.

Through YourEncore, you can contract with a retiree who has relevant experience for a specific, short-term assignment (compensation is based on the person's preretirement salary, adjusted for inflation). For example, we might tap a former Boeing engineer with expertise in virtual aircraft design to apply his or her skills in virtual product prototyping and manufacturing design at P&G, even though our projects have nothing to do with aviation. What makes this model so powerful is that client companies can experiment at low cost and with little risk on cross-disciplinary approaches to problem solving. At any point, we might have 20 retirees from YourEncore working on P&G problems.

Yet2.com. Six years ago, P&G joined a group of *Fortune* 100 companies as an initial investor in Yet2.com, an online marketplace for intellectual property exchange. Unlike NineSigma and InnoCentive, which focus on helping companies find solutions to technology problems, Yet2.com brokers technology transfer both into and out of companies, universities, and government labs. Yet2.com works with clients to write briefs describing the technology that they're seeking or making available for license or purchase, and distributes these briefs throughout a global network of businesses, labs, and institutions. Network members interested in posted briefs contact Yet2.com and request an introduction to the relevant client. Once introduced, the parties negotiate directly with each other. Through Yet2.com, P&G was able to license its low-cost microneedle technology to a company specializing in drug delivery. As a result of this relationship, we have ourselves licensed technology that has applications in some of our core businesses.

When to Engage

Once products and ideas are identified by our networks around the world, we need to screen them internally. All the screening methods are driven by a core understanding, pushed down through the entire organization, of what we're looking for. It's beyond the scope of this article to describe all of the processes we use to evaluate ideas from outside. But a look at how we might screen a new product found by a technology entrepreneur illustrates one common approach.

When our technology entrepreneurs are meeting with lab heads, scanning patents, or selecting products off store shelves, they're conducting an initial screening in real time: Which products, technologies, or ideas meet P&G's where-to-play criteria? Let's assume a technology entrepreneur finds a promising product on a store shelf that passes this initial screening. His or her next step will be to log the product into our online "eureka catalog," using a template that helps organize certain facts about the product: What is it? How does it meet our business needs? Are its patents available? What are its current sales? The catalog's descriptions and pictures (which have a kind of Sharper Image feel) are distributed to general managers, brand managers, R&D teams, and others throughout the company worldwide, according to their interests, for evaluation.

Meanwhile, the technology entrepreneur may actively promote the product to specific managers in relevant lines of business. If an item captures the attention of, say, the director of the baby care business, she will assess its alignment with the goals of the business and subject it to a battery of practical questions–such as whether P&G has the technical infrastructure needed to develop the product – meant to identify any showstopping impediments to development. The director will also gauge the product's business potential. If the item continues to look promising, it may be tested in consumer panels and, if the

The Osaka Connection

In the connect-and-develop world, chance favors the prepared mind. When one of P&G's technology entrepreneurs discovered a stain-removing sponge in a market in Osaka, Japan, he sent it to the company for evaluation. The resulting product, the Mr. Clean Magic Eraser, is now in third-generation development and has achieved double its projected revenues.

German chemical company BASF manufactures a melamine resin foam called Basotect for soundproofing and insulation in the construction and automotive industries.

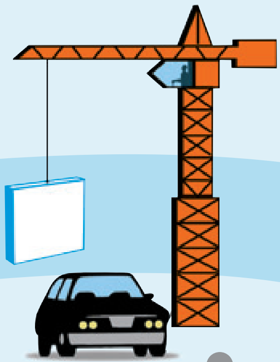

LEC, a Tokyo-based consumer-products company, markets Basotect foam in Japan as a household sponge called Cleenpro.

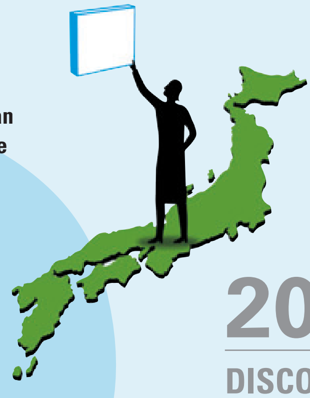

2001

DISCOVER

A Japan-based technology entrepreneur with P&G discovers the product in an Osaka grocery store, evaluates its market performance in Japan, and establishes its fit with the P&G home-care product development and marketing criteria.

2002

EVALUATE

The technology entrepreneur sends samples to R&D product researchers in Cincinnati for performance evaluation and posts a product description and evaluation of market potential on P&G's internal "eureka catalog" network.

Market research confirms enthusiasm for the product. The product is moved into portfolio for development; P&G negotiates purchase of Basotect from BASF and terms for further collaboration.

2003

LAUNCH

Basotect is packaged as-is and launched nationally as Mr. Clean Magic Eraser.

Mr. Clean Magic Eraser is launched in Europe.

BASF and P&G researchers collaborate in shared labs to improve Basotect's cleaning properties, durability, and versatility.

2004

COCREATE

The first cocreated Basotect product, the Magic Eraser Duo, is launched nationally in the United States.

The cocreated Magic Eraser Wheel & Tire is launched nationally in the United States.

BASF and P&G continue to collaborate on next-generation Magic Eraser products.

response is positive, moved into our product development portfolio. Then we'll engage our external business development (EBD) group to contact the product's manufacturer and begin negotiating licensing, collaboration, or other deal structures. (The EBD group is also responsible for licensing P&G's intellectual property to third parties. Often, we find that the most profitable arrangements are ones where we both license to and license from the same company.) At this point, the product found on the outside has entered a development pipeline similar in many ways to that for any product developed in-house.

The process, of course, is more complex and rigorous than this thumbnail sketch suggests. In the end, for every 100 ideas found on the outside, only one ends up in the market.

Push the Culture

No amount of idea hunting on the outside will pay off if, internally, the organization isn't behind the program. Once an idea gets into the development pipeline, it needs R&D, manufacturing, market research, marketing, and other functions pulling for it. But, as you know, until very recently, P&G was deeply centralized and internally focused. For connect and develop to work, we've had to nurture an internal culture change while developing systems for making connections. And that has involved not only opening the company's floodgates to ideas from the outside but actively promoting internal idea exchanges as well.

For any product development program, we tell R&D staff that they should start by finding out whether related work is being done elsewhere in the company; then they should see if an external source – a partner or supplier, for instance – has a solution. Only if those two avenues yield nothing should we consider inventing a solution from scratch. Wherever the solution comes from (inside or out), if the end product succeeds in the marketplace, the rewards for employees involved in its development are the same. In fact, to the extent that employees get recognition for the speed of product development, our reward systems actually favor innovations developed from outside ideas since, like Pringles Prints, these often move more quickly from concept to market.

We have two broad goals for this reward structure. One is to make sure that the best ideas, wherever they come from, rise to the surface. The other is to exert steady pressure on the culture, to continue to shift mind-sets away from resistance to "not invented here." Early on, employees were anxious that connect and develop might eliminate jobs or that P&G would lose capabilities. That stands to reason, since as you increase the ideas coming in from the outside you might expect an equivalent decrease in the need for internal ideas. But with our growth objectives, there is no limit to our need for solid business-building ideas. Connect and develop has not eliminated R&D jobs, and it has actually required the company to develop new skills. There are still pockets within P&G that have not embraced connect and develop, but the trend has been toward accepting the approach, even championing it, as its benefits have accrued and people have seen that it reinforces their own work.

Words of Warning

Procter & Gamble's development and implementation of connect and develop has unfolded over many years. There have been some hiccups along the way, but largely it has been a methodical process of learning by doing, abandoning what doesn't work and expanding what does. Over five years in, we've identified three core requirements for a successful connect-and-develop strategy.

• Never assume that "ready to go" ideas found outside are truly ready to go. There will always be development work to do, including risky scale-up.

• Don't underestimate the internal resources required. You'll need a full-time, senior executive to run any connect-and-develop initiative.

• Never launch without a mandate from the CEO. Connect and develop cannot succeed if it's cordoned off in R&D. It must be a top-down, companywide strategy.

Adapt or Die

We believe that connect and develop will become the dominant innovation model in the twenty-first century. For most companies, as we've argued, the alternative invent-it-ourselves model is a sure path to diminishing returns.

To succeed, connect and develop must be driven by the top leaders in the organization. It is destined to fail if it is seen as solely an R&D strategy or isolated as an experiment in some other corner of the company. As Lafley did at P&G, the CEO of any organization must make it an explicit company strategy and priority to capture a certain amount of innovation externally. In our case, the target is a demanding – even radical – 50%, but we're well on our way to achieving it.

Don't postpone crafting a connect-and-develop strategy, and don't approach the process incrementally. Companies that fail to adapt to this model won't survive the competition.

Reprint R0603C; HBR OnPoint 351X
To order, see page 151.

MANAGEMENT SCIENCE
Vol. 53, No. 7, July 2007, pp. 1113–1126
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5307 | 1113
DOI 10.1287/mnsc.1060.0624

Interfirm Collaboration Networks: The Impact of Large-Scale Network Structure on Firm Innovation

Melissa A. Schilling
Stern School of Business, New York University, 40 West Fourth Street, New York, New York 10012, mschilli@stern.nyu.edu

Corey C. Phelps
Department of Management and Organization, University of Washington, Box 353200, Seattle, Washington 98195, cphelps@u.washington.edu

The structure of alliance networks influences their potential for knowledge creation. Dense local clustering provides information transmission capacity in the network by fostering communication and cooperation. Nonredundant connections contract the distance between firms and give the network greater reach by tapping a wider range of knowledge resources. We propose that firms embedded in alliance networks that exhibit both high clustering and high reach (short average path lengths to a wide range of firms) will have greater innovative output than firms in networks that do not exhibit these characteristics. We find support for this proposition in a longitudinal study of the patent performance of 1,106 firms in 11 industry-level alliance networks.

Key words: alliances; networks; innovation; patents
History: Accepted by Brian Uzzi and Luis Amaral, special issue editors; received June 16, 2004. This paper was with the authors 5 months for 2 revisions.

Introduction

Although research has long recognized the importance of interfirm networks in firm innovation (see Freeman 1991 for a review), much of this work has treated the network concept as a metaphor. Only recently have researchers begun to assess the formal structural properties of alliance networks and their impact on firm innovation. This research has focused on a firm's position within a broader network of relationships or the structure of its immediate network neighborhood rather than the structure of the overall network. Studies have examined a firm's centrality (Smith-Doerr et al. 1999), number of alliances (Shan et al. 1994), and local network structure (Ahuja 2000, Baum et al. 2000). To our knowledge, empirical research has not yet examined the impact of the structure of industry-level[1] alliance networks on member firm innovation. In a related study, however, Uzzi and Spiro (2005) examined the network structure of the creative artists who made broadway musicals from 1945 to 1989, and concluded that the large-scale structure of the artists' collaboration network significantly influenced their creativity, and the financial and artistic performance of their musicals. This raises the following questions: Does the structure of an industry-level interfirm network influence the rate of knowledge creation among firms in the network? If so, what structural properties will enhance firm innovation?

To address these questions, we examine the impact of two key large-scale network properties, clustering and reach, on the innovative output of members of the network. The dense connectivity of clusters creates *transmission capacity* in a network (Burt 2001), enabling large amounts of information to rapidly diffuse, while reach (i.e., short path lengths to a wide range of firms) ensures that diverse information sources can be tapped. We argue that networks with both high clustering and high reach will significantly enhance the creative output of member firms. We test this hypothesis using longitudinal data on the innovative performance of a large panel of firms operating in 11 industry-level alliance networks.

This research offers several important contributions for understanding knowledge creation in interfirm networks. First, we find empirical support for our argument that the combination of clustering and reach increases member firm innovation. To our knowledge, no other study has attempted to assess the effect of industry-level interfirm networks on the innovation performance of member firms. Although recent studies have examined the structure of large-scale interfirm networks and the possible causes of these structures (Baum et al. 2003, Kogut and Walker

[1] An industry-level network is a specific type of whole or "large-scale" network. Wellman (1988, p. 26) defined a whole network as the relationships that exist among members of a population.

2001), little research has examined the *consequences* of large-scale network structure in an industrial setting (Uzzi and Spiro 2005 is a recent exception). Second, while most studies of network structure have examined a single industry, our study uses longitudinal data on 11 industries, which strengthens the generalizability of our findings.

We begin by describing two key structural characteristics of interfirm networks and their effect on information diffusion in the network. From this we develop a hypothesis about how the structure of interfirm networks will influence the innovative output of member firms. We test the hypothesis on a large, unbalanced panel of firms embedded in 11 industry-level alliance networks.

Large-Scale Interfirm Networks and Firm Knowledge Creation

We adopt a recombinatory search perspective in explaining the process of innovation (Fleming 2001). Innovation is characterized as a problem-solving process in which solutions to problems are discovered via search (Dosi 1988). Prior research suggests that search processes that lead to the creation of new knowledge, embodied in artifacts such as patents and new products, most often involve the novel recombination of known elements of knowledge, problems, or solutions (Fleming 2001, Nelson and Winter 1982) or the reconfiguration of the ways in which knowledge elements are linked (Henderson and Clark 1990). Critical inputs into this process include access to and familiarity with a variety of knowledge elements (e.g., different technological components and the scientific and engineering know-how embedded in them), novel problems and insights into their resolution, failed recombination efforts, and successful solutions (Hargadon and Fanelli 2002). Firms that have greater access to and understanding of these recombinatory resources should be advantaged in their innovation efforts.

As firms form and maintain alliances with each other, they weave a network of direct and indirect relationships. As a result, firms embedded in these networks gain access to information and know-how of direct partners and that of others in the network to which they are indirectly connected (Ahuja 2000, Gulati and Gargiulo 1999). The network of alliance relationships constitutes a conduit that channels the flow of information and know-how among firms in the network (Ahuja 2000, Owen-Smith and Powell 2004), with each member firm acting as both a recipient and transmitter of information (Ahuja 2000). The structure of these networks greatly influences the dynamics of information diffusion within the networks. Large-sample studies have found that *direct* alliance relationships facilitate knowledge flows between partners (Gomes-Casseres et al. 2006, Mowery et al. 1996) and enhance the innovative performance of firms (e.g., Deeds and Hill 1996, Stuart 2000). Research also shows that the extent to which a firm is *indirectly* connected to other firms in an alliance network enhances its innovativeness (Ahuja 2000, Owen-Smith and Powell 2004, Soh 2003).

Given the role of direct and indirect ties as channels for the flow of information and know-how, we argue that the structure of the interfirm network will significantly influence the recombination process. Two structural characteristics that have a particularly important role in diffusion are *clustering* and *reach*.

Clustering

Alliance networks tend to be highly clustered: Some groups of firms will have more links connecting them to each other than to the other firms in the network. A firm's clustering coefficient can be calculated as the proportion of its partners that are themselves directly linked to each other. The clustering coefficient of the overall network is the average of this measure across all firms in the network. Several mechanisms lead to clustering in interfirm knowledge networks, but two of the most common are linking based on similarity or complementarity. Firms tend to interact more intensely or frequently with other firms with which they share some type of proximity or similarity, such as geography or technology (Baum et al. 2003, Rosenkopf and Almeida 2003). This tends to result in a high degree of clustering.

Clustering increases the information transmission capacity of a network. First, the dense connectivity of individual clusters ensures that information introduced into a cluster will quickly reach other firms in the cluster. The multiple pathways between firms also enhance the fidelity of the information received. Firms can compare the information received from multiple partners, helping them to identify ways in which it has been distorted or is incomplete. Second, clusters within networks are important structures for making information exchange meaningful and useful. The internal density of a cluster can increase the dissemination of alternative interpretations of problems and their potential solutions, deepening the collective's understanding and stimulating collective problem solving (Powell and Smith-Doerr 1994). The development of a shared understanding of problems and solutions greatly facilitates communication and further learning (Brown and Duguid 1991, Powell et al. 1996). Third, dense clustering can make firms more willing and able to exchange information (Ahuja 2000). Sociologists (e.g., Coleman 1988, Granovetter 1992) have suggested that densely clustered networks give rise to trust, reciprocity norms, and a shared identity, all of which lead to a high level of cooperation and can facilitate collaboration by providing self-enforcing informal governance mechanisms (Dyer and Singh 1998). In addition to stimulating greater transparency, trust and reciprocity exchanges facilitate intense interaction among personnel from partnered firms (Uzzi 1997), improving the transfer of tacit, embedded knowledge (Hansen 1999, Zander and Kogut 1995). Thus, clustering enables richer and greater amounts of information and knowledge to be exchanged and integrated more readily.

When dense clusters are sparsely connected to each other, they become important structures for creating and preserving the requisite variety of knowledge in the broader network that enables knowledge creation. The internal cohesion of a cluster can cause much of the information and knowledge shared within a cluster to become homogeneous and redundant (Burt 1992, Granovetter 1973). The dense links provide many redundant paths to the same actors, and thus the same sources of information and knowledge. Cohesion can also lead to norms of adhering to established standards and conventions, which can potentially stifle experimentation and creativity (Uzzi and Spiro 2005). This limits innovation. Clusters of firms will, however, tend to be heterogeneous across a network in the knowledge they possess and produce due to the different initial conditions and causes for each cluster to form. The diversity of knowledge distributed across clusters in the network provides the requisite variety for recombination.

Clustering thus offers both local and global advantages. Firms benefit from having redundant connectivity among their immediate neighbors because it enhances the speed and likelihood of information access, and the depth of information interpretation. Firms also benefit from being embedded within a larger network that is clustered because the information a firm receives from partners that are embedded in other clusters is likely to be more complete and richly understood than information received from partners not embedded in clusters, and because information received from different clusters is likely to be diverse, enabling a wider range of recombinatorial possibilities.

Reach

The size of a network and its average path length (i.e., the average number of links that separates each pair of firms in the network) also impacts information diffusion and novel recombination. The more firms that can be reached by any path from a given firm, the more knowledge that firm can potentially access. However, the likelihood, speed, and integrity of knowledge transfer between two firms are directly related to the path length separating those two firms. The diffusion of information and knowledge occurs more rapidly and with more integrity in networks with short average path lengths than in networks with longer paths (Watts 1999). A firm that is connected to a large number of firms by a short average path can reach more information, and can do so quickly and with less risk of information distortion than a firm that is connected to fewer firms or by longer paths. To capture this we use *distance-weighted reach*.

A firm's distance-weighted reach is the sum of the reciprocal distances to every firm that is reachable from a given firm, i.e., $\sum_j 1/d_{ij}$, where d_{ij} is defined as the minimum distance (geodesic), d, from a focal firm i to partner j, where $i \neq j$. A network's *average distance-weighted reach* is this measure averaged across all firms in the network, $(\sum_n \sum_j 1/d_{ij})/n$, where n is the number of firms in the network. Other things being equal, a very large connected network with a very short average path length (e.g., a completely connected network where there are many firms and every firm is directly connected to every other firm, or a star graph with many firms all connected to the same central "hub" firm) will have the greatest average distance-weighted reach. Longer path lengths, smaller network size, or disconnects that fragment the network into multiple components all decrease average distance-weighted reach.

The preceding reveals one of the key benefits of using distance-weighted reach: It provides a meaningful measure of the overall size and connectivity of a network, even when that network has multiple components, and/or component structure is changing over time. It avoids the infinite path length problem typically associated with disconnected networks by measuring only the path length between connected pairs of nodes, and it provides a more meaningful measure than the simple average path length between connected pairs by factoring in the size of connected components.[2]

Because forming alliances is costly and constrained, there appears to be a trade-off between forming dense clusters to facilitate rapid exchange and integration of knowledge, versus forging links to create short paths to a wider range of firms. However, recent research has shown that even sparse, highly clustered networks can have high reach if there are a few links creating bridges between clusters (Watts 1999, Hansen 2002, Hargadon 1998). Bridges between clusters of firms provide member firms access to diverse information that exists beyond their local cluster, enabling

[2] The authors are grateful to Steve Borgatti for pointing this out. They are also grateful to Mark Newman for numerous discussions about how to handle the infinite path length consideration in our networks.

new combinations with their existing knowledge sets, while preserving the information transmission advantages of clusters. As Uzzi and Spiro (2005) note, bridges between clusters increase the likelihood that different ideas and routines will come into contact, enabling recombinations that incorporate both previous conventions and novel approaches. The combination of clustering and reach thus enables a wide range of information to be exchanged and integrated rapidly, leading to greater knowledge creation. In sum, we predict a multiplicative interaction between clustering and reach in their effect on firm knowledge creation. Consistent with the symmetrical nature of such interactions (Jaccard and Turrisi 2003), we have argued and expect that the effect of clustering on firm knowledge creation will be increasingly positive as reach increases, while the effect of reach on knowledge creation will be increasingly positive as clustering increases.

Hypothesis. *Firms participating in alliance networks that combine a high degree of clustering and reach will exhibit more knowledge creation than firms in networks that do not exhibit these characteristics.*

Methods

To test our hypothesis, we constructed a large, unbalanced panel of U.S. firms for the period 1990–2000. The panel includes all U.S. firms that were part of the alliance networks of 11 high-technology manufacturing industries: aerospace equipment (standard industrial classifications (SICs)): 3721, 3724, 3728, 3761, 3764, 3769; automotive bodies and parts (3711, 3713, 3714); chemicals (281-, 282-, 285-, 286-, 287-, 288-, 289-); computer and office equipment (3571, 3572, 3575, 3577); household audiovisual equipment (3651); medical equipment (3841, 3842, 3843, 3844, 3845); petroleum refining and products (2911, 2951, 2952, 2992, 2999); pharmaceuticals (2833, 2834, 2835, 2836); semiconductors (3674); telecommunications equipment (366-), and measuring and controlling devices (382-).

The choice of industries was particularly important for this study. The 11 industries selected have been designated as high technology in numerous Bureau of Labor Statistics studies (e.g., Hecker 1999).[3] These industries provide an excellent context for our study for three reasons. First, knowledge creation is fundamental to the pursuit of competitive advantage in high-technology industries. Second, firms in these industries actively use alliances in pursuit of their innovation activities (Vonortas 1997). Third, because we use patent data for our dependent variable, it is important to select industries that use patents. There is evidence that firms in these industries actively patent their inventions (Levin et al. 1987).

Alliance Networks

We chose to measure the network structure created by publicly reported strategic alliances for two reasons. First, there is a rich history of research on the importance of strategic alliances as a mechanism for knowledge sharing among firms (Freeman 1991, Gulati 1998, Powell et al. 1996). Second, alliances are used by a wide range of firms (both public and private) in a wide range of industries, and are often used explicitly for the exchange and joint creation of knowledge.

Social network research has identified three procedural tactics for establishing network boundaries for empirical research: attributes of actors that rely on membership criteria, such as membership in an industry; types of relations between actors, such as participation in strategic alliances; and participation in a set of common events (Laumann et al. 1983). Accordingly, we employed two rules to guide our construction of the 11 industry networks used in this study. First, each alliance included at least one firm that was a member of the target industry (indicated by its primary four-digit SIC). Second, each alliance had to operate in the target industry, as indicated by its primary four-digit SIC of activity.

Alliance data were gathered using Thomson Corp.'s SDC Platinum database. The SDC data have been used in a number of empirical studies on strategic alliances (e.g., Anand and Khanna 2000, Sampson 2004). For each industry, alliances were collected that were announced between 1990 and 1997. We chose 1990 as the initial year for our sample because information on alliances formed prior to 1990 is very sparse in the SDC database (Anand and Khanna 2000, p. 300). Separate alliance networks were created for each industry according to the alliance's primary SIC. Both public and private firms were included. We use data on only U.S. firms because the SDC alliance data are much more complete for U.S. firms than for non-U.S. firms (Phelps 2003). All alliances were aggregated to the parent corporation.

The resulting data set includes 1,106 firms involved in 3,663 alliances. Many of the alliances included more than two participating firms, so the number of dyads is greater, totaling 5,306. Because any type of alliance may provide a path for knowledge diffusion, and because prior studies indicate that the breadth of an alliance's true activity is often much greater than what is formally reported (Powell et al. 1996), we include all alliance types in our analysis. We do, however,

[3] We omitted high-tech manufacturing industries that rarely use alliances: special industry machinery (355), electrical industrial apparatus (362), search and navigation equipment (381), and photographic equipment and supplies (386).

control for the proportion of alliances in each network formed for the explicit purpose of technology exchange or development.

Alliances typically last for more than one year, but alliance termination dates are rarely reported. This required us to make an assumption about alliance duration. We took a conservative approach and assumed that alliance relationships last for three years, consistent with recent empirical work on average alliance duration (Phelps 2003). Other research has taken a similar approach, using windows ranging from one to five years (e.g., Gulati and Gargiulo 1999, Stuart 2000). We created alliance networks based on three-year windows (i.e., 1990–1992, 1991–1993, ... 1995–1997), resulting in six snapshots of network structure for each industry, for a total of 66 alliance network snapshots. Each network snapshot was constructed as an undirected binary adjacency matrix (Wasserman and Faust 1994).[4] Multiple alliances between the same pair of firms in a time window were treated as one link. UCINET 6 was used to obtain measures on these networks, as described below (Borgatti et al. 2002).

As we focus on publicly reported contractual alliance agreements, we do not observe the numerous informal collaborative arrangements that exist between firms in our sample. Such informal arrangements often lead to the types of formal agreements that we observe (Powell et al. 1996, Rosenkopf et al. 2001). Thus, our analysis represents a conservative test of our diffusion argument because our data do not include informal relationships that promote knowledge transfer.

Dependent Variable: Patents

One way that knowledge creation is instantiated is in the form of inventions (Schmookler 1966). Knowledge embedded in artifacts such as inventions represents the "empirical knowledge" of organizations (Hargadon and Fanelli 2002). Inventions thus provide a trace of an organization's knowledge creation. Patents provide a measure of novel invention that is externally validated through the patent examination process (Griliches 1990). Patent counts have been shown to correlate well with new product introductions and invention counts (Basberg 1987). Trajtenberg (1987) concluded that patents are valid and robust indicators of knowledge creation. One of the challenges with using patents to measure innovation is that the propensity to patent may vary with industry, resulting in a potential source of bias (Levin et al. 1987). We addressed this potential bias in three ways. First, we sample only high-tech manufacturing industries, which helps to ensure a degree of commonality in the industries' emphasis on innovation. To further capture differences in emphasis on innovation, we control for industry-level R&D intensity. Third, to control for unobserved factors that influence the propensity to patent that are likely to be stable within industries, we control for industry fixed effects. The propensity to patent may also differ due to firm characteristics (Griliches 1990). We attempt to control for such sources of heterogeneity using covariate, *Presample Patents* (described below), and firm fixed and random effects in our estimations.

We measure the dependent variable, $Patents_{it}$, as the number of successful patent applications for firm i in year t. We used the Delphion database to collect yearly patent counts for each of the firms, aggregating subsidiary patents up to the ultimate parent level. Granted patents were counted in their year of application. Yearly patent counts were created for each firm for the period of 1993 to 2000, enabling us to assess different lag specifications between alliance network structure and patent output.

Independent Variables

Clustering Coefficient. To measure the clustering in each network for each time period, we used the weighted overall clustering coefficient measure (Borgatti et al. 2002, Newman et al. 2002):

$$Clustering_w = \frac{3 \times (\text{number of triangles in the graph})}{(\text{number of connected triples})},$$

where a *triangle* is a set of three nodes (e.g., i, j, k), each of which is connected to both of the others, and a *connected triple* is a set of three nodes in which at least one is connected to both the others (e.g., i is connected to j and k, but j and k need not be connected). This measure indicates the proportion of triples for which transitivity holds (i.e., if i is connected to j and k, then by transitivity, j and k are connected). The factor of three in the numerator ensures that the measure lies strictly in the range of zero and one because each triangle implies three connected triples. The weighted overall clustering coefficient represents the percentage of a firm's alliance partners that are also partnered with each other, weighted by the number of each

[4] Each matrix reflects the alliances maintained within the network as of the end of the focal year. Because alliances often endure longer than one year, constructing adjacency matrices using only alliances announced in the focal year would bias the connectivity of the observed networks downward. Consider the initial year of the panel for the network variables (1992): Using only alliances formed in 1992 would not capture the alliance relationships formed prior to, yet maintained through, 1992. Data on both presample alliance formation and alliance duration is needed to accurately assess network structure in each of the sample years. Moving three-year windows more accurately reflects the structure of an alliance network in the annual adjacency matrices. Robinson and Stuart (2007) use a similar approach in assessing alliance networks in the biotechnology industry.

firm's partners, averaged across all firms in the network. This variable can range from zero to one, with larger values indicating increasing clustering. While network density captures the density of the entire network, the clustering coefficient captures the degree to which the overall network contains localized pockets of dense connectivity. A network can be globally sparse and still have a high clustering coefficient.

Reach. To capture the reach of each network for each time period, we use a measure of average distance-weighted reach (Borgatti et al. 2002). This is a compound measure that takes into account both the number of firms that can be reached by any path from a given firm, and the path length it takes to reach them. This measure is calculated as

$$\textit{Average distance weighted reach} = \left[\sum_{n}\sum_{j} 1/d_{ij}\right] \Big/ n,$$

where n is the number of nodes in the network, and d_{ij} is defined as the minimum distance (geodesic), d, from a focal node i to partner j, where $i \neq j$. Average distance-weighted reach can range from 0–n, with larger values indicating higher reach.

Clustering × Reach. We predict that the combination of clustering and reach will have a positive impact on member firm innovation, and thus include the interaction term, *Clustering × Reach*.

Firm-Level Control Variables

Presample Patents. To control for unobserved heterogeneity in firm patenting, we follow the presample information approach of Blundell et al. (1995) and calculate the variable *Presample Patents* as the sum of patents obtained by a firm in the five years prior to its entry into the sample.

Betweenness Centrality. We control for the possibility that firms that occupy more central positions in alliance networks may generate more innovations than more peripheral firms (e.g., Owen-Smith and Powell 2004, Soh 2003). We operationalize *Centrality* using Freeman's (1979) measure of "betweenness centrality," which captures the extent to which a firm is located on the shortest path (i.e., geodesic) between any two actors in its alliance network. Formally, betweenness centrality for firm i in year t is calculated as

$$\textit{Betweenness Centrality}_{it} = \sum_{j<k} g_{jk}(n_i)/g_{jk},$$

where $g_{jk}(n_i)$ refers to the number (n) of geodesics (i.e., shortest paths) linking firms j and k that contain focal firm i. The term $g_{jk}(n_i)/g_{jk}$ captures the probability that firm i is involved in the shortest path between j and k. Betweenness centrality is the sum of these estimated probabilities over all pairs of firms (excluding the ith firm) in the network. We use normalized betweenness centrality (i.e., betweenness divided by maximum possible betweenness, expressed as a percentage) to make the measure comparable across time and industry networks.

Local Efficiency. While studies have found that the extent to which a firm's partners are nonredundant enhances its knowledge creation (Baum et al. 2000), other research shows that redundant links improve innovation (Ahuja 2000). Although the empirical evidence is mixed, controlling for the effect of local structural holes is important if we wish to demonstrate that the global structure of the alliance network in which a firm is embedded has an independent and significant influence on its subsequent patenting. We control for the influence of a firm's local network structure using Burt's (1992) measure of efficiency. Efficiency captures the extent to which a firm's partners are nonredundant, indicating the presence of structural holes. Local efficiency for firm i in year t is computed as

$$\textit{Local Efficiency}_{it} = \left\lfloor \sum_{j} \left\lfloor 1 - \sum_{q} p_{iq} m_{iq} \right\rfloor \right\rfloor \Big/ N_i, \quad j \neq q,$$

where p_{iq} is the proportion of i's relations invested in the relationship with q, m_{jq} represents the marginal strength of the relationship between alter j and alter q (as we use binary data, values of m_{jq} are set to one if the relationship is present and zero otherwise), and N_i represents the number of unique alliance partners connected to firm i. This measure can range from zero to one, with higher values indicating greater efficiency.

Industry (Network) Control Variables

Network Density. We control for the overall density of the network with the variable *Network Density*, calculated for each industry network and time period. We do so because the rate and extent to which information diffuses increases with density (Yamaguchi 1994). This variable measures the ratio of existing links in the network to the number of possible pairwise combinations of firms, and may range from zero to one, with larger values indicating increasing density.

Centralization. The extent to which a network is centralized can also influence its diffusion properties. A highly centralized network is one in which all ties run through one or a few nodes, thus decreasing the distance between any pair of nodes (Wasserman and Faust 1994). To control for network centralization, we employ Freeman's (1979) index of group betweenness centralization, calculated for each industry network and time period. Group betweenness centralization for network j in year t is

$$\textit{Betweenness Centralization}_{jt} = 100 \times \left\{ \sum_{i=1}^{g} [C'_B(n^*) - C'_B(n_i)]/(g-1) \right\},$$

where $C'_B(n^*)$ is the largest realized normalized betweenness centrality for the set of firms in network j in year t, $C'_B(n_i)$ is the normalized betweenness centrality for firm i (in industry network j for year t), and g is the number of firms. This variable is expressed as a percentage and can range from zero, where all firms have the same individual betweenness centrality, to 100, where one firm connects all other firms.

Industry R&D Intensity. To control for differences in the emphasis on and costliness of innovation across industries, we employ a time-varying measure of industry-level R&D intensity (R&D expenditures/sales). We collected annual R&D expenditures and sales of firms in each industry from Compustat. We would have preferred to control for R&D intensity at the firm level; however, nearly 42% of our sample firms were privately owned during some portion of the sample, and R&D expenditures are not available for private firms. In investigating the robustness of our results, we utilize a control variable (stock of patents obtained in the past four years) that has been shown to be highly correlated with annual firm-level R&D expenses. Our results are unchanged when including this variable in our models.

Proportion of Alliances for R&D, Cross-Technology Transfer, or Licensing. Alliances that are established for the purpose of technology exchange or development may be more directly related to firm patenting. To examine this possibility, we include a time-varying measure of the percentage of alliance agreements in each network that were established explicitly for the purpose of joint research and development, cross-technology transfer, or technology licensing.

Model Specification

The dependent variable in this study, *Patents*, is a count variable and takes on only nonnegative integer values. The linear regression model is inadequate for modeling such variables because the distribution of residuals will be heteroscedastic nonnormal. A Poisson regression approach is appropriate to model count data (Hausman et al. 1984). However, the Poisson distribution contains the strong assumption that the mean and variance are equal. Patent data often exhibit overdispersion, where the variance exceeds the mean (Hausman et al. 1984). In the presence of overdispersion, coefficients will be estimated consistently, but their standard errors will generally be underestimated, leading to spuriously high levels of significance (Cameron and Trivedi 1986). Each model that we report, when estimated using the Poisson specification, exhibited significant overdispersion.

A commonly used alternative to the Poisson regression model is the negative binomial model. The negative binomial model is a generalization of the Poisson model and allows for overdispersion by incorporating an individual, unobserved effect into the conditional mean (Hausman et al. 1984). The panel data implementation of the negative binomial model accommodates explicit control of persistent individual unobserved effects through both fixed and random effects. In the present study, unobserved heterogeneity refers to the possibility that unmeasured (or unmeasurable) differences among observationally equivalent firms affects their patenting. Unobserved heterogeneity may also stem from unmeasured, systematic time period and industry effects. Failing to control for such unobserved heterogeneity, if present, can result in specification error (Heckman 1979).

We employ a number of strategies to control for these sources of unobserved heterogeneity. First, we include year fixed effects to control for systematic period effects such as differences in macroeconomic conditions that may affect all sampled firms' patent rates. Second, we employ individual firm effects to control for unobserved, temporally stable firm differences in patenting. We use both firm fixed and random effects in alternative estimations of our model. We use the Hausman et al. (1984) implementation of fixed effects in the context of a negative binomial model, which employs a conditional maximum-likelihood estimation procedure.[5]

[5] Allison and Waterman (2002) recently criticized the Hausman et al. (1984) conditional negative binomial fixed-effects model as not being a "true" fixed-effects method in that it does not control for all time-invariant covariates. Allison and Waterman (2002) developed an unconditional negative binomial model that uses dummy variables to represent fixed effects, which effectively controls for all stable individual effects. This procedure has been implemented in Limdep 8.0. However, estimates of β are inconsistent in negative binomial models when using such a dummy variable approach in short panels, due to the incidental parameters problem (Cameron and Trivedi 1998, p. 282). The number of unit-specific (e.g., firm) parameters (α_i) increases with the sample size, while the number of periods (T) stays fixed, resulting in a limited number of observations to estimate a large number of parameters. Contrary to linear regression models, the maximum-likelihood estimates for α_i and β are not independent for negative binomial models because the inconsistency of the estimates of α_i are transmitted into the maximum likelihood estimate of β. Given that this method is a true fixed-effects specification, it does not allow for time-invariant covariates such as *Presample Patents*. Thus, we chose not to employ Allison and Waterman's (2002) unconditional estimator. We report the results using the Hausman et al. (1984) conditional fixed-effects

We also use Hausman et al.'s random effects specification, which assumes that overdispersion due to unobserved heterogeneity is randomly distributed across firms. Because the random effects specification assumes that the unobserved firm-specific effect is uncorrelated with the regressors, we report the results from both fixed and random effects as a robustness check.

As an additional control for firm-level unobserved heterogeneity, we adopt the presample information approach of Blundell et al. (1995). Blundell et al. (1995) argued that because the main source of unobserved heterogeneity in models of innovation lies in the different knowledge stocks with which firms enter a sample, a variable that approximates the build-up of firm knowledge at the time of entering the sample is a particularly good control for unobserved heterogeneity. The *Presample Patents* variable described above serves as a control for unobserved differences in firm knowledge stocks upon their entry into the sample. Blundell et al. (1995) showed that the use of a presample patent entry stock measure virtually eliminated persistent serial correlation in their panel data models. We also include industry dummies in our models to control for unobserved industry effects that are not captured by the firm effects.

A final estimation issue concerns the appropriate lag structure of the independent variables. Based on prior research that investigates the relationship between interfirm alliances and innovation (e.g., Ahuja 2000, Sampson 2004, Stuart 2000), we employ alternative lags of our independent variables relative to our dependent variable. We estimate models using one-year, two-year, and three-year lags. We do so to explore the robustness of our findings across alternative specifications. All models were estimated with Limdep 8.0. The model we estimate takes the general form provided below (aerospace is the omitted industry and 1992 is the omitted year). Variables are indexed across firms (i), industry (j), and time (t):

$$
\begin{aligned}
&\mathit{Patents}_{it+1(2,3)} \\
&\quad = f(\mathit{Clustering}_{jt}, \mathit{Reach}_{jt}, \mathit{Clustering} * \mathit{Reach}_{jt}, \\
&\qquad \mathit{R\&D\ Alliance\ \%}_{jt}, \mathit{R\&D\ Intensity}_{jt}, \mathit{Centrality}_{it}, \\
&\qquad \mathit{Local\ Efficiency}_{it}, \mathit{Centralization}_{jt}, \mathit{Density}_{jt}, \\
&\qquad \mathit{Presample\_Patents}_{it}, \mathit{Automotive}, \mathit{Chemicals}, \\
&\qquad \mathit{Computers}, \mathit{Audiovisual}, \mathit{Medical}, \mathit{Petroleum}, \\
&\qquad \mathit{Pharmaceuticals}, \mathit{Semiconductors}, \\
&\qquad \mathit{Telecommunications}, \mathit{Measuring}, 1993, 1994, 1995, \\
&\qquad 1996, 1997).
\end{aligned}
$$

approach. We point out that the results we obtained from both fixed- and random effects specifications are highly consistent (see the Results section). Studies that have employed both the Hausman et al. (1984) negative binomial fixed-effects approach and that of Allison and Waterman (2002) have found very similar results (e.g., Dee et al. 2005).

Results

A summary of the network size and component structure for each industry, averaged over time, is provided in Table 1. As shown, there is substantial variation across industries in the number of firms that participate in alliances. This is largely due to differences in industry size. The average number of alliances per firm within each industry exhibits much less variation. The next column provides the average number of firms in each network. This number includes firms from the industry and their partners, some of which are not in the target industry. The next column indicates the percentage of nodes in the network that are connected to the single largest ("main") component. This number varies significantly both across industry and over time (not shown). While researchers often study only the main component, in our study this would have yielded misleading results. Whereas in some industries there is a large main component that is relatively stable over time (e.g., pharmaceuticals), in other industries there are multiple large components, and those components merge and split apart over time. For example, between 1996 and 1997 in the computer industry, a large component broke away from the main component (see Figure 1). If we had focused only on the single largest component, we would have both understated the amount of

Table 1 **Network Size and Component Structure, Averages over 1992–2000**

Industry	Average number of firms from industry in alliances[a]	Average number of alliances per firm	Average network size (nodes)[b]	Percent in main component (%)
Aerospace	9	3.05	28	46
Automotive	15.67	3.43	53.2	37
Chemicals	45.17	2.97	199.8	11
Computers and office equipment	79.67	4.48	347	45
Household audiovisual equipment	9	1.5	28.3	10
Measuring and controlling	22.67	1.96	48.33	21
Medical equipment	66.17	1.66	172.33	7
Petroleum refining and products	5.3	2.65	24.83	18
Pharmaceuticals	218.33	2.54	510	64
Semiconductors	58.67	3.51	204	55
Telecommunication equipment	44.83	6.53	266.33	54

[a]This number includes only those firms with the designated primary SICs; it does not include partners in the network that are not in those SICs.

[b]Includes all U.S. firms in the network, including both those with the designated primary SICs and their alters, regardless of SIC.

Figure 1 Network Size and Component Structure (Common Shade of Gray Indicates Firms in Same Component)

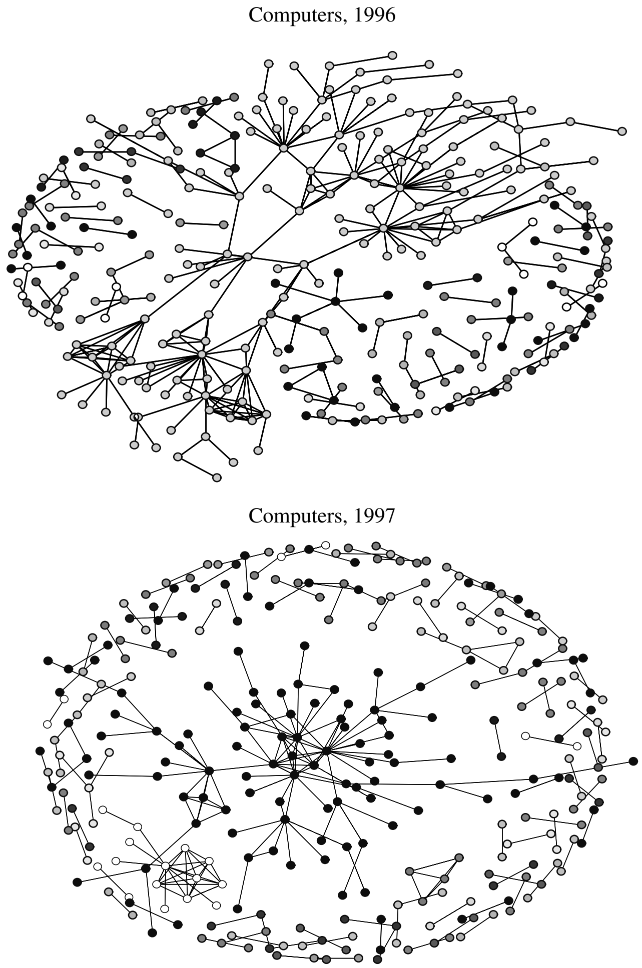

alliance activity in the industries, and overstated the amount of change in alliance activity over time.

Table 2 reports the negative binomial panel regression results for the three dependent variables ($Patents_{it+1}$; $Patents_{it+2}$; $Patents_{it+3}$). Because the random effects specification assumes that regressors and firm-specific effects are uncorrelated, we also provide results using firm fixed effects as a robustness check. Separate results are provided for three dependent variables. Models 1, 2, and 3 report the results using a one-year lag between the independent variables and firm patenting ($Patents_{it+1}$). Models 4, 5, and 6 report the results using a two-year lag ($Patents_{it+2}$), and Models 7, 8, and 9 report the results using a three-year lag ($Patents_{it+3}$). For each dependent variable, the first models (1, 4, and 7) include the control variables only, the second models add the direct effects of *Clustering* and *Reach* (Models 2, 5, and 8), and the third model adds the interaction term, *Clustering* × *Reach* (Models 3, 6, and 9). To conserve space, firm, industry, and time period effects, while estimated, are not reported.

Our sole hypothesis predicted a positive effect of the interaction of *Clustering* and *Reach* on firm patenting. The interaction term, *Clustering* × *Reach*, does not obtain statistical significance at conventional levels in the model specified with a one-year lag, using either fixed or random firm effects (Model 3). The coefficient for *Clustering* × *Reach* is positive and statistically significant in models using both two- and three-year lags (Models 6 and 9). This result holds for models using both fixed and random firm effects. Thus, our hypothesis received strong support in models using two- and three-year lags.[6]

To better understand the meaning of the interaction effect, the nature of the coefficients for *Clustering* and *Reach* in Models 6 and 9 in Table 2 must be understood. The estimated coefficients for *Clustering* and *Reach* in these models are *simple* effects rather than true main effects due to the significance of the interaction term (Jaccard and Turrisi 2003). Consequently, the effect of each on *Patents* is conditioned on the other variable taking on the value of zero. For example, the coefficient estimate of −0.022 for *Reach* in Model 6 (Random Effects) assumes that the value of *Clustering* is equal to zero (thus removing the impact of the interaction with *Reach*). Thus, the negative sign on the coefficient for *Reach* cannot be interpreted as a negative (main) effect of *Reach* on *Patents*. While the effect of *Reach* is indeed negative when *Clustering* is zero, the effect becomes positive when values of *Clustering* exceed 0.267 (the range of *Clustering* in the data is 0.0–0.8). Similarly, the effect of *Clustering* is negative (although not statistically significant) when *Reach* is equal to zero, but becomes positive for values of *Reach* greater than 1.224 (the range of *Reach* is 1.88–61.18).[7] The fact that the effects of both *Clustering* and *Reach* become positive when the other obtains a relatively small value, and increase in their positive effects with increases in the other, provides further support for our hypothesis. These mutually reinforcing effects are consistent with the symmetrical nature of multiplicative interaction effects (Jaccard and Turrisi 2003).

Plots of the effect of the interaction on predicted values of $Patents_{t+2}$ and $Patents_{t+3}$ reinforce this interpretation. For ease of presentation and interpretation, we used the log-linear form of the negative binomial models in Table 2 (i.e., where the log of

[6] We also conducted a test of the hypothesis at the industry level rather than the firm level. In this test, we regressed the industry's average number of firm patents on the network- and industry-level variables. We obtained nearly identical results to those in Table 2. These results are available from the authors upon request.

[7] To calculate these effects, we used the log-linear form of the negative binomial models in Table 2 (i.e., where the log of the conditional mean function is linear in the estimated parameters). We followed the approach for calculating interaction effects described by Jaccard and Turrisi (2003, p. 23).

Table 2 Panel Negative Binomial Regression Models with Fixed and Random Effects ($N = 1{,}106$; $Obs = 3{,}444$)

	$Patents_{it+1}$			$Patents_{it+2}$			$Patents_{it+3}$		
	1	2	3	4	5	6	7	8	9
Fixed effects									
Constant	1.136** (0.354)	0.582 (0.359)	0.604 (0.360)	1.257** (0.327)	1.663** (0.333)	1.614** (0.324)	1.433** (0.337)	1.859** (0.369)	1.825** (0.368)
Presample Patents	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)
Density	−0.248 (1.154)	−0.624 (1.358)	−0.527 (1.468)	−0.411 (1.529)	−2.220 (1.808)	−2.637 (1.843)	−2.012 (1.861)	−1.598 (2.509)	−1.674 (2.134)
Centralization	−0.014 (0.008)	−0.014 (0.008)	−0.012 (0.008)	−0.018** (0.006)	−0.016* (0.007)	−0.035** (0.006)	0.019** (0.007)	0.019** (0.007)	0.019* (0.007)
Ind. R&D Intensity	2.739 (2.668)	2.867 (2.522)	2.877 (2.581)	0.741 (2.366)	−0.088 (2.373)	−0.246 (2.327)	−7.126** (2.478)	−6.754** (2.504)	−6.754** (2.504)
R&D Alliance %	−0.112 (0.275)	0.223 (0.275)	0.222 (0.289)	0.068 (0.217)	−0.131 (0.223)	−0.188 (0.191)	−0.040 (0.248)	−0.305 (0.264)	−0.312 (0.304)
Efficiency	−0.199** (0.068)	−0.189** (0.072)	−0.190** (0.073)	−0.303** (0.091)	−0.321** (0.095)	−0.327** (0.087)	−0.267** (0.097)	−0.272** (0.089)	−0.270** (0.088)
Betweenness	0.003 (0.006)	0.003 (0.005)	0.003 (0.005)	0.005 (0.006)	0.004 (0.007)	0.002 (0.006)	−0.001 (0.009)	−0.001 (0.010)	−0.001 (0.010)
Clustering		0.420** (0.136)	0.507* (0.235)		0.346** (0.127)	−0.141 (0.196)		0.234 (0.183)	−0.319 (0.279)
Reach		0.010** (0.003)	0.011** (0.003)		−0.012** (0.003)	−0.020** (0.004)		−0.007* (0.003)	−0.009* (0.004)
Clustering × Reach			−0.015 (0.030)			0.081** (0.023)			0.014* (0.007)
Log Likelihood	−4,646.65	−4,637.32	−4,637.12	−4,597.46	−4,586.78	−4,577.98	−4,468.75	−4,464.64	4,464.46
Random effects									
Constant	1.118** (0.309)	0.542 (0.339)	0.541 (0.339)	0.984** (0.307)	1.342** (0.303)	1.256** (0.290)	0.920** (0.296)	1.333** (0.331)	1.214** (0.321)
Presample Patents	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)	0.001** (0.000)
Density	1.444 (0.900)	0.250 (1.092)	0.243 (1.166)	0.527 (1.197)	−1.872 (1.394)	−2.451 (1.352)	−1.454 (1.434)	−1.286 (1.618)	−1.538 (1.654)
Centralization	−0.021** (0.006)	−0.020** (0.007)	−0.021** (0.007)	−0.021** (0.006)	−0.020** (0.006)	−0.027** (0.005)	0.016* (0.006)	0.017* (0.007)	0.013* (0.006)
Ind. R&D Intensity	0.887 (2.429)	1.030 (2.408)	1.027 (2.424)	−0.357 (2.231)	−0.818 (2.151)	−0.590 (2.135)	−8.029** (2.278)	−7.987** (2.343)	−8.101** (2.460)
R&D Alliance %	0.014 (0.230)	0.383 (0.214)	0.384 (0.222)	0.208 (0.215)	−0.017 (0.187)	−0.090 (0.158)	0.106 (0.220)	−0.139 (0.233)	−0.153 (0.274)
Efficiency	−0.342** (0.062)	−0.336** (0.069)	−0.336** (0.069)	−0.396** (0.079)	−0.436** (0.081)	−0.435** (0.073)	−0.297** (0.087)	−0.307** (0.080)	−0.312** (0.078)
Betweenness	0.008 (0.005)	0.007 (0.004)	0.007 (0.005)	0.003 (0.005)	0.004 (0.005)	0.001 (0.005)	−0.000 (0.008)	−0.001 (0.008)	−0.001 (0.008)
Clustering		0.554** (0.106)	0.548** (0.212)		0.485** (0.116)	−0.101 (0.186)		0.152 (0.159)	−0.422 (0.344)
Reach		0.008** (0.003)	0.008* (0.003)		−0.013** (0.003)	−0.022** (0.003)		−0.008* (0.003)	−0.011** (0.004)
Clustering × Reach			0.001 (0.028)			0.082** (0.019)			0.043* (0.020)
a	0.707** (0.047)	0.716** (0.047)	0.710** (0.048)	0.675** (0.047)	0.684** (0.048)	0.690** (0.480)	0.650** (0.046)	0.652** (0.046)	0.652** (0.046)
b	0.358** (0.021)	0.360** (0.022)	0.360** (0.022)	0.321** (0.019)	0.328** (0.020)	0.334** (0.02)	0.291** (0.018)	0.290** (0.018)	0.293** (0.018)
Log likelihood	−8,520.70	−8,509.78	−8,509.78	−8,425.33	−8,407.95	−8,392.95	−8,198.66	8,194.98	−8,193.03

Notes. All models include firm, time period, and industry effects. Standard errors are in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$ (two-tailed tests for all variables).

the conditional mean function is linear in the estimated parameters) to calculate these effects. Figure 2 presents the interaction plot of *Clustering* and *Reach* to illustrate the magnitude of the interaction effect. The "Low Clustering" line shows the slope of the effect of *Reach* on *Patents* when the value of *Clustering* is set to one standard deviation below its mean. The end points of the line are calculated at one standard deviation below and above the mean of *Reach*. The "High Clustering" line represents the effect of *Reach* on *Patents* when the value of *Clustering* is set to one standard deviation above its mean. Consistent with the results in Models 6 and 9 of Table 2, increases in *Reach* increase the positive effect of *Clustering* on *Patents*. The symmetrical case of plotting low and high *Reach* lines for low and high values of *Clustering* (not shown) provides similar results.

To assess the magnitude of the interaction effect we employed the estimated marginal effects ($e^{\beta X}\beta$). The magnitude of the interaction effect when both component variables increase one standard deviation above their means for the model employing a two-year lag and random effects is 1.00 patents (or 2.3%). For the model specified with a two-year lag and employing firm fixed effects, this yielded an increase of 0.98 patents (for the average firm), or 2.3%. The magnitude of the interaction effect is smaller in the models using a three-year lag. Thus, the size of the interaction effect in absolute terms is fairly small in our data and appears to realize its peak within two years. Based on these results, we speculate that the effect of network structure as a medium of knowledge diffusion decays over time. While a particular structure may persist over time, the knowledge that diffuses through it has limited benefit as actors absorb and apply these knowledge flows to productive ends.

The results related to the control variables also merit discussion. The effect of betweenness centrality on subsequent firm patenting failed to achieve statistical significance in any of the estimated models. In contrast, efficiency had a significant negative effect on firm patenting in all models. This result suggests that the presence of structural holes in a firm's ego network of alliance relationships has a deleterious effect on its inventive output. This is consistent with results obtained by Ahuja (2000) and Soh (2003). To our knowledge, our study represents the largest panel data investigation of this relationship.

Among the other variables in the models, most were not consistent in terms of sign and significance. This might be due, in part, to the moderate-to-large correlations among the network measures (i.e., *Centralization*, *Density*, *Reach*, *Clustering*, and *Clustering* × *Reach*). This multicollinearity might influence the robustness of our main finding because parameter estimates are unstable to very small changes in the data when substantial collinearity is present, sometimes resulting in the signs on estimated coefficients

Figure 2 **Graph of Interactions for Random Effects Models,** $Patents_{t+2}$ **and** $Patents_{t+3}$

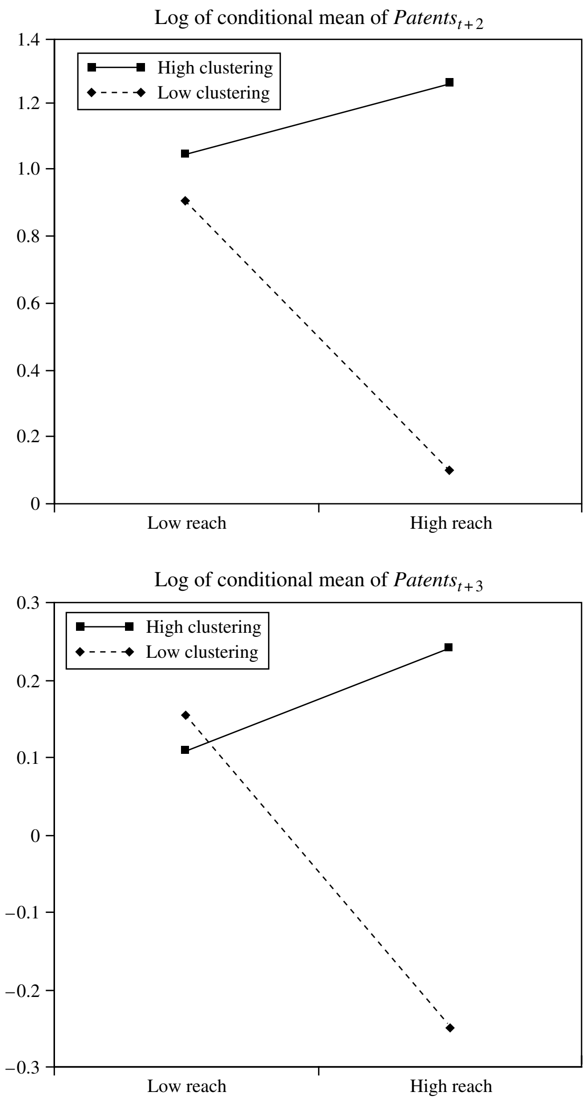

to flip (known as the "wrong sign" problem) (Gujarati 1995). To examine the influence of multicollinearity on our main result, we reran each of the models in Table 2 with *Centralization* removed and, alternatively, with *Density* removed (not reported here). The results for *Reach*, *Clustering*, and *Clustering* × *Reach* remained substantively unchanged across all models.

Finally, the *Presample Patents* variable was positive and significant in all models, indicating its importance as a control for firm-level unobserved heterogeneity. Furthermore, several time period and industry dummies (not reported) were consistently significant in all models.

Robustness of Results

One concern regarding our results is that we were not able to control for differences in firm R&D because nearly 42% of our sample firms were privately owned during some portion of the sample. Prior research has found that patent stock measures and annual R&D expenditures are highly correlated (e.g., Trajtenberg 1990). We measured a firm's patent stock as the total number of patents obtained by firm i in the four years prior to and including year t. Due to the extremely high correlation between this variable and *Presample Patents* ($r = 0.94$), we reestimated all of our models using the time-varying patent stock variable in place of *Presample Patents*. As might be expected (due to the substantial correlation between the two variables) our results (not reported) did not substantively change from those reported in Table 2.

For our second robustness check, we analyzed the data using a Poisson fixed-effects estimation procedure. We did so to address the concern identified in Footnote 5. This approach controls for all unobserved time-invariant sources of heterogeneity. In this analysis, we *excluded* all time-invariant variables and obtained qualitatively similar results (not reported) to those presented in Table 2.

A third concern regarding our results is that they may be influenced by the presence of persistent serial correlation in the residuals. This could result from temporally stable unobserved firm effects (Greene 1997), or from reverse causality running from firm invention to industry-level network structure (e.g., clustering or reach), manifesting in the lagged network variables. We explicitly address the first potential source of serial correlation by including firm fixed effects. Unreliable estimates may also result from unobservables that vary systematically over time. Serial correlation in the errors would persist even after controlling for stable firm effects. We examined this possibility in two ways. First, we regressed our measures of clustering and reach, and their interaction on annual firm patent counts using a linear panel data model. We did so using contemporaneously measured firm patents and one-, two-, and three-year lags of firm patents. We found no significant relationship between firm patents and clustering, reach, or their interaction in any of these models. Next, we aggregated firm patents to the industry level using the average annual patent count across firms in the industry. The idea here is that as industry inventiveness increases, so does the likelihood that firms in such industries form alliance networks with high clustering and reach. We ran the same specifications as those using firm patents and found no significant effects.

Discussion

We argued that two key structural properties of large-scale networks, clustering and reach, play important roles in network diffusion and search. Clustering enables even a globally sparse network to achieve high information transmission capacity through locally dense pockets of closely connected firms. Reach increases the quantity and diversity of information available to firms in the network by bringing the information resources of more firms within relatively close range. We thus argued that networks that have both the high information transmission capacity enabled by clustering, and the high quantity and diversity of information provided by reach, should facilitate greater innovation by firms that are members of the network. We tested this argument using longitudinal data on the innovative performance of a large panel of firms operating in 11 industry-level alliance networks. The results indicated support for our argument: the combination of clustering and reach was associated with significantly higher firm patenting. The results were stronger for models employing a two- and three-year lag versus a one-year lag, suggesting that firms do not quickly realize the innovation benefits of collaboration (Stuart 2000). These results were robust to a number of controls and model specifications.

Our results support much of the theory developed in recent work on small-world networks (Cowan and Jonard 2003, Uzzi and Spiro 2005). Our results are consistent with Uzzi and Spiro's argument that the cohesion and connectivity of a small-world network enable the circulation of creative material that can be recombined into new creative products. Our argument that the heterogeneity of knowledge distributed across clusters enhances innovation is similar to Uzzi and Spiro's argument that the different conventions and styles used in different clusters are a valuable source of diversity in the network.[8]

This research has a number of contributions. First, whereas previous alliance network research has examined the impact of a firm's network position or the structure of its immediate network neighborhood on firm innovation, our study is the first that we know of to examine the influence of the structure of industry-level alliance networks on firm innovation. The results of this study also inform the debate over whether innovation is enhanced by network density or efficiency (see Ahuja 2000): Both *local* density and *global* efficiency can exist simultaneously, and it is this combination that enhances innovation.

Finally, our results speak to the literature on knowledge spillovers. Knowledge spillovers represent an externality in which the knowledge produced by one firm can be appropriated, at little cost, by other firms (Jaffe 1986). Empirical evidence indicates that spillovers are important in explaining innovation and productivity growth (Griliches 1992). However, spillovers are not equally accessible to or appropriable by all firms. Prior research has shown that spillovers tend to be spatially bounded: Their effect is more pronounced for firms conducting research in similar technological domains (Jaffe 1986) and geographic locations (Feldman 1999). Our results add to this literature by suggesting that interfirm networks may be an important mechanism of knowledge spillovers, and that the specific pattern these relationships exhibit can have important consequences for the innovativeness of networked firms.

We do not wish to overstate our results—this study has a number of limitations. Our findings may be influenced by our assumption of average alliance duration. If alliances endure, on average, for more than three years, then the connectivity of our observed networks will be biased downward. This bias may influence our results. Unfortunately, due to data limitations, we were unable to explore this possibility.[9] A limitation of our theoretical focus is that we ignore the influence of network characteristics other than structure. We do not address the properties of the alliances themselves (e.g., strength, governance structure, scope). Different types of relationships may be better or worse for searching for, versus transferring, knowledge (Hansen 1999). In addition, different types of relationships will be more or less costly to maintain, and thus affect the efficiency of network structure for knowledge creation. We do not examine how

[8] Uzzi and Spiro's data and analysis are different from ours in some important ways. First, as they point out (2005, p. 470, Footnote 8), in a mature small-world network such as theirs, the path length changes little over time, behaving like a fixed effect with a constant value near one. This means that that their principle finding is driven primarily by temporal variation in clustering. Our networks, by contrast, exhibit significant cross-sectional and temporal variation in path length and network size, leading to great variation in our measure of reach. Second, our networks are far less dense than their network. Their network becomes sufficiently dense and clustered that it leads to excessive cohesion and homogenization of material, and a decline in creative performance. In essence, such a globally dense network has the advantages and disadvantages we argued would exist within each cluster. To investigate this effect in our data, we reestimated each of our models, replacing our interaction term with the quadratic version of clustering (i.e., clustering^2). This variable was not significant in any model; thus, we have no evidence of a parabolic effect of clustering in our data. We speculate that our networks never reach a high level of density, and thus are at less risk of excessive cohesion. Finally, and perhaps most importantly, Uzzi and Spiro's network is composed of individuals, whereas our networks are composed of firms. Some of the dynamics that lead to deleterious effects of cohesion (for example, strong feelings of obligation between friends leading to an "assistance club" for ineffectual members of the network) are far more likely in the relationships between individuals than between firms.

[9] We did not collect alliance formation data prior to 1990 because SDC data prior to that time is inconsistent. We chose to end observations of patents in 2000 (implying our alliance observation ended in 1997) because the lag between patent application and grant date is two to four years (which was toward the end of our data collection).

the attributes of the firms shape the flow of knowledge (see Owen-Smith and Powell 2004). We have also not addressed the potential impact of the nature of knowledge that is being accessed, transferred, and recombined in the network. Different characteristics of knowledge (e.g., tacit versus explicit, complex versus simple, etc.) can influence the knowledge creation and innovation process (Zander and Kogut 1995). Network structure may also differentially interact with different dimensions of knowledge. For example, the high density of clusters may facilitate the search and transfer of tacit, complex knowledge, but the relatively few connections to other clusters may make such search and transfer problematic. These aspects of relationships and knowledge will likely be important in fully understanding the relationship between interfirm knowledge networks and knowledge creation, but are beyond the scope of our paper.

Another limitation of our work is that the generalizability of our main result is likely to be limited to industries that make frequent use of alliances. Networks characterized by extreme sparsity may not have a sufficient degree of connectedness to observe clustering or meaningful reach. However, the implications of our results are not necessarily limited to alliance relationships. Because firms are connected via other relationships, the global structure of such relationships may influence firm innovativeness. For example, firms are often connected by interpersonal collaborative relationships among individual inventors. The extent to which the global structure of these relationships is characterized by clustering and reach may have implications for the inventiveness of individual inventors and their firms (Fleming et al. 2004). Furthermore, because knowledge can flow between firms through other mechanisms such as individual mobility, geographic clustering, participation in technical committees, or learning from information made public through patenting, it is possible that some of the knowledge creation advantages of a particular alliance network structure might spill over to other industry (or nonindustry) participants. Each of these limitations represents an exciting area for future research.

Acknowledgments
This research was supported by the National Science Foundation under Grant SES-0234075. The authors are grateful for the suggestions of Juan Alcácer, Laszlo Barabasi, Joel Baum, Bill Greene, Anne Marie Knott, Dan Levinthal, Bill McKelvey, Mark Newman, Joe Porac, Lori Rosenkopf, Rob Salomon, Kevin Steensma, Kate Stovel, and Duncan Watts.

References
Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Admin. Sci. Quart.* **45** 425–455.

Allison, P. D., R. Waterman. 2002. Fixed-effects negative binomial regression models. *Soc. Methodology* **32** 247–265.

Anand, B. N., T. Khanna. 2000. Do firms learn to create value? The case of alliances. *Strategic Management J.* **21** 295–315.

Basberg, B. L. 1987. Patents and the measurement of technological change: A survey of the literature. *Res. Policy* **16** 131–141.

Baum, J. A. C., T. Calabrese, B. S. Silverman. 2000. Don't go it alone: Alliance network composition and startups' performance in Canadian biotechnology. *Strategic Management J.* **21** 267–294.

Baum, J. A. C., A. V. Shipilov, T. J. Rowley. 2003. Where do small worlds come from? *Indust. Corporate Change* **12** 697–725.

Blundell, R. R., R. Griffith, J. Van Reenen. 1995. Dynamic count data models of technological innovation. *Econom. J.* **105** 333–344.

Borgatti, S. P., M. G. Everett, L. C. Freeman. 2002. Ucinet for Windows: Software for social network analysis. Harvard, Analytic Technologies.

Brown, J., P. Duguid. 1991. Organizational learning and communities of practice: Towards a unified view of working, learning and innovation. *Organ. Sci.* **2** 40–57.

Burt, R. S. 1992. *Structural Holes*. Harvard University Press, Cambridge, MA.

Burt, R. S. 2001. Bandwidth and echo: Trust, information, and gossip in social networks. A. Casella, J. E. Rauch, eds. *Networks and Markets: Contributions from Economics and Sociology*. Russel Sage Foundation, New York, 30–74.

Cameron, A. C., P. K. Trivedi. 1986. Econometric models based on count data: Comparisons and applications of some estimators and tests. *J. Appl. Econometrics* **1** 29–53.

Cameron, A. C., P. K. Trivedi. 1998. *Regression Analysis of Count Data*. Cambridge University Press, Cambridge, UK.

Coleman, J. S. 1988. Social capital in the creation of human capital. *Amer. J. Sociol.* **94** S95–S120.

Cowan, R., N. Jonard. 2003. The dynamics of collective invention. *J. Econom. Behav. Organ.* **52** 513–532.

Dee, T. S., D. C. Grabowski, M. A. Morrisey. 2005. Graduated driver licensing and teen traffic fatalities. *J. Health Econom.* **24** 571–589.

Deeds, D. L., C. W. L. Hill. 1996. Strategic alliances and the rate of new product development: An empirical study of entrepreneurial firms. *J. Bus. Venturing* **11** 41–55.

Dosi, G. 1988. Sources, procedures, and microeconomic effects of innovation. *J. Econom. Literature* **26** 1120–1171.

Dyer, J. H., H. Singh. 1998. The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Acad. Management Rev.* **23** 660–679.

Feldman, M. P. 1999. The new economics of innovation, spillovers, and agglomeration: A review of empirical studies. *Econom. Innovation New Tech.* **8** 5–25.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47** 117–132.

Fleming, L., A. Juda, C. King, III. 2004. Small worlds and regional innovative advantage. Harvard Business School Working Paper Series, 04-008, Harvard University, Cambridge, MA.

Freeman, C. 1991. Networks of innovators: A synthesis of research issues. *Res. Policy* **20** 499–514.

Freeman, L. 1979. A set of measures of centrality: I. Conceptual clarification. *Soc. Networks* **1** 215–239.

Gomes-Casseres, B., J. Hagedoorn, A. Jaffe. 2006. Do alliances promote knowledge flows? *J. Financial Econom.* **80** 5–33.

Granovetter, M. S. 1973. The strength of weak ties. *Amer. J. Sociol.* **78** 1360–1380.

Granovetter, M. S. 1992. Problems of explanation in economic sociology. N. Nohria, R. Eccles, eds. *Networks and Organizations: Structure, Form, and Action*. Harvard Business School Press, Boston, MA, 25–56.

Greene, W. H. 1997. *Econometric Analysis*, 3rd ed. Prentice Hall, Upper Saddle River, NJ.

Griliches, Z. 1990. Patent statistics as economic indicators: A survey. *J. Econom. Literature* **28** 1661–1707.

Griliches, Z. 1992. The search for R&D spillovers. *Scandinavian J. Econom.* **94**(Suppl.) 29–47.

Gujarati, D. N. 1995. *Basic Econometrics*, 3rd ed. McGraw-Hill, New York.

Gulati, R. 1998. Alliances and networks. *Strategic Management J.* **19** 293–317.

Gulati, R., M. Gargiulo. 1999. Where do interorganizational networks come from? *Amer. J. Sociol.* **104** 1439–1493.

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.

Hansen, M. T. 2002. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Organ. Sci.* **13** 232–250.

Hargadon, A. B. 1998. Firms as knowledge brokers: Lessons in pursuing continuous innovation. *California Management Rev.* **40**(3) 209–227.

Hargadon, A. B., A. Fanelli. 2002. Action and possibility: Reconciling dual perspectives of knowledge in organizations. *Organ. Sci.* **13** 290–302.

Hausman, J., B. Hall, Z. Griliches. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica* **52** 909–938.

Hecker, D. 1999. High technology employment: A broader view. *Monthly Labor Rev.* **122**(6) 18–28.

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

Henderson, R., K. Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* **35** 9–30.

Jaccard, J., R. Turrisi. 2003. *Interaction Effects in Multiple Regression*. Sage, Thousand Oaks, CA.

Jaffe, A. 1986. Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits and market value. *Amer. Econom. Rev.* **76** 984–1001.

Kogut, B. 2000. The network as knowledge. Generative rules and the emergence of structure. *Strategic Management J.* **21** 405–425.

Kogut, B., G. Walker. 2001. The small world of Germany and the durability of national networks. *Amer. Sociol. Rev.* **66** 317–335.

Laumann, E. O, P. V. Marsden, D. Prensky. 1983. The boundary specification problem in network analysis. R. S. Burt, M. J. Minor, eds. *Applied Network Analysis*. Sage, Beverly Hills, CA, 18–34.

Levin, R., A. Klevorick, R. Nelson, S. Winter. 1987. Appropriating the returns from industrial research and development. *Brookings Papers Econom. Activity, Microeconomics* **3** 783–820.

Mowery, D. C., J. E. Oxley, B. S. Silverman. 1996. Strategic alliances and interfirm knowledge transfer. *Strategic Management J.* **17** 77–91.

Nelson, R. R., S. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.

Newman, M. E. J., S. H. Strogatz, D. J. Watts. 2002. Random graph models of social networks. *Proc. Natl. Acad. Sci. USA* **99** 2566–2572.

Owen-Smith, J., W. W. Powell. 2004. Knowledge networks as channels and conduits: The effects of spillovers in the Boston biotechnology community. *Organ. Sci.* **15** 5–21.

Phelps, C. 2003. Technological exploration: A longitudinal study of the role of recombinatory search and social capital in alliance networks. Unpublished dissertation, New York University, New York.

Powell, W. W., L. Smith-Doerr. 1994. Networks and economic life. N. J. Smelser, R. Swedberg, eds. *The Handbook of Economic Sociology*. Princeton University Press, Princeton, NJ, 368–402.

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** 116–145.

Robinson, D. T., T. E. Stuart. 2007. Network effects in the governance of strategic alliances. *J. Law, Econom., Organ.* **23**(1) 242–273.

Rosenkopf, L., P. Almeida. 2003. Overcoming local search through alliances and mobility. *Management Sci.* **49** 751–766.

Rosenkopf, L., A. Metiu, V. P. George. 2001. From the bottom up? Technical committee activity and alliance formation. *Admin. Sci. Quart.* **46** 748–772.

Sampson, R. 2004. The cost of misaligned governance in R&D alliances. *J. Law, Econom., Organ.* **20** 484–526.

Schmookler, J. 1966. *Invention and Economic Growth*. Harvard University Press, Cambridge, MA.

Shan, W., G. Walker, B. Kogut. 1994. Interfirm cooperation and startup innovation in the biotechnology industry. *Strategic Management J.* **15** 387–394.

Smith-Doerr, L., J. Owen-Smith, K. W. Koput, W. W. Powell. 1999. Networks and knowledge production: Collaboration and patenting in biotechnology. R. Leenders, S. Gabbay, eds. *Corporate Social Capital*. Kluwer Academic Publishers, Norwell, MA, 331–350.

Soh, P.-H. 2003. The role of networking alliances in information acquisition and its implication for new product performance. *J. Bus. Venturing* **18** 727–744.

Stuart, T. E. 2000. Interorganizational alliances and the performance of firms: A study of growth and innovation rates in a high technology industry. *Strategic Management J.* **21** 791–812.

Trajtenberg, M. 1987. Patents, citations, and innovations: Tracing the links. Working Paper 2457, National Bureau of Economic Research, Cambridge, MA.

Trajtenberg, M. 1990. A penny for your quotes: Patent citations and the value of innovations. *RAND J. Econom.* **21** 172–187.

Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. *Admin. Sci. Quart.* **42** 35–67.

Uzzi, B., J. Spiro. 2005. Collaboration and creativity: The small world problem. *Amer. J. Sociol.* **111** 447–504.

Vonortas, N. S. 1997. Research joint ventures in the US. *Res. Policy* **26** 577–595.

Wasserman, S., K. Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, UK.

Watts, D. J. 1999. Networks, dynamics, and the small-world phenomenon. *Amer. J. Sociol.* **105** 493–528.

Wellman, B. 1988. Structural analysis: From method and metaphor to theory and substance. B. Wellman, S. D. Berkowitz, eds. *Social Structures: A Network Approach*. Cambridge University Press, Cambridge, UK, 19–61.

Yamaguchi, K. 1994. The flow of information through social networks: Diagonal-free measures of inefficiency and the structural determinants of inefficiency. *Soc. Networks* **16** 57–86.

Zander, U., B. Kogut. 1995. Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test. *Organ. Sci.* **6** 76–92.

© Academy of Management Journal
2004, Vol. 47, No. 6, 876–892.

PREFERENTIAL LINKAGE AND NETWORK EVOLUTION: A CONCEPTUAL MODEL AND EMPIRICAL TEST IN THE U.S. VIDEO GAME SECTOR

N. VENKATRAMAN
CHI-HYON LEE
Boston University

We examined how network structure (density overlap and embeddedness) and technology characteristics of a platform (dominance and newness) shaped interorganizational coordination of product launches in the U.S. video game industry. We found that the developers' choices to launch games for particular game consoles were significantly explained by these four factors using multiprobability regression on a primary data set of 2,815 launches between 1995 and 2002. This analysis was complemented with application of a network visualization technique.

During the last two decades, there has been a marked increase in research attention to interorganizational relationships that reflect organizational economics and sociological perspectives. Empirical studies have focused on motives for the formation and governance of these relationships in different settings (e.g., Gulati, 1995) and on the benefits of such relationships (e.g., Ahuja, 2000). These studies have led to strong calls for recognition of a relational view of strategy in which "[a] firm's critical resources may span firm boundaries and may be embedded in interfirm routines and processes" (Dyer & Singh, 1998: 681). Further, scholars have called for increased efforts to understand management choices and actions taking a network perspective that includes a focus on a firm's set of relationships with suppliers, buyers, and "complementors," or other firms providing coordinating products (Gulati, Nohria, & Zaheer, 2000; Katz & Shapiro, 1994).

Most studies with a network perspective have focused on vertical relationships involving buyers and/or suppliers along the value chain. Researchers have invoked transaction cost economics (Williamson, 1991) to test predictions about the formation and governance of interorganizational vertical relationships (e.g., Heide, 1994; Subramani & Venkatraman, 2003). Few studies have, however, focused on relationships involving two companies who function as *complementors* by providing coordinated products that jointly appeal to end customers (Baldwin & Clark, 2000; Garud, Kumaraswamy, & Langlois, 2003). This type of relationship is particularly important in high-technology settings, where a particular standard may fail to become dominant because support from a complementor network is inadequate (Katz & Shapiro, 1994; Schilling, 2003). IBM personal computers became dominant in the 1980s via the support of third-party developers of compatible software, and the dominance of Microsoft's Windows platform is largely due to support from complementary software developers (Bresnahan & Greenstein, 1999; McKenzie, 2000).

Understanding the network of complementor relationships is important for several reasons. First, these relationships confer resources critical to the success of many contemporary high-technology ecosystems comprising coordinated products such as hardware and software, video game consoles and games, mobile phones and applications, and so on (Frels, Shervani, & Shrivastava, 2003; Shapiro & Varian, 1998). Second, the relationships between complementors involve coordinated product launches and mutual dependency for success but no typical mechanisms of long-term contracts or equity investments (Gawer & Cusumano, 2002). Three, these relationships provide insights into the path-dependency of technological evolution (Arthur, Durlauf, & Lane, 1997).

We thank the Boston University School of Management's Institute for Leading in a Dynamic Economy (BUILDE) for providing support for completion of this study. Discussions with Bala Iyer helped sharpen some of the ideas presented here. Byunggon Yoo assisted in data collection. Early versions of this research were presented at seminars at the New York University Stern School of Management, the University of Minnesota Carlson School of Management, the London Business School, an Academy of Management symposium, Seattle, and the Marshall School of Management at the University of Southern California. Joe Galaskiewicz, Akbar Zaheer, and the *AMJ* reviewers provided valuable suggestions for restructuring and repositioning the paper. Persephone Doliner provided invaluable editorial support in developing the manuscript. All limitations are our own.

Our study is distinctive in the following ways: we focused on the dynamic network of relationships between the manufacturers of video game consoles (also termed "video game platforms") and video game developers over an eight-year period. In particular, we sought to explain video game developers' decisions to launch new games for specific platforms on the basis of the characteristics, at the time of each decision, of the network structure and specific platform features. These decisions affect the success of both the developers and their complementors, namely the platform manufacturers, since their success is interdependent. Our approach to explaining developers' decisions to link to different platforms recognizes the time-varying nature of network structure and platform characteristics. These decisions collectively drive the evolution of the network (for an analogous approach in the context of choices of friends over time, see Van de Bunt, Van Duijn & Snijders, 1999).

Our study is an effort to understand how "networks evolve and change over time" (Nohria, 1992: 15) in the U.S. video game sector. Thus, our work is in the spirit of recent attempts to discern logics of network evolution; for example, Baum, Shipilov, and Rowley (2003) focused on how "small world structures" emerge and evolve over time by testing scenarios for the formation of ties in investment bank syndicate networks in Canada; and Powell, White, Koput, and Owen-Smith (2004) modeled the evolution of the biotechnology field in terms of how multiple actors—commercial entities, universities, research institutes, venture capitalists, and small start-ups—interacted over time.

PREFERENTIAL LINKS IN NETWORK FORMATION

What explains a developer's decision to launch a new game title for a particular platform at a given time? We recognize the inherent competition for links among developers striving to dynamically balance their sets of links across their sets of platforms. Game developers reveal their relationship intent with every game launch: they can magnify the importance of a platform by making the game exclusive on it; they can "port" a game from one platform to another and reduce the differentiation between them. Moreover, game developers constantly strive to have their games on the most popular platform while seeking to avoid highly competitive segments in which their games may not be differentiated and platforms that may be at the declining stage of the technology life cycle. Thus, the launch of new titles reveals developer intentions that influence the nature of relationships with platform manufacturers.

We derive our hypotheses on preferential links by integrating various theoretical strands from population ecology, social network theory, and technological evolution. Two hypotheses deal with macro network structure, namely, how a platform's "density overlap" and "embeddedness" with developers influence platform choice. The other two address a platform's intrinsic characteristics, namely, its market dominance and newness. The strength of our model is its integrated focus on a broader set of characteristics as influences on a firm's choice about network links.

Density Overlap between Platforms as an Influence on Preferential Linkage

Within population ecology, the density dependence model focuses on the dynamics of competition (organizations undermining each other's fates) and mutualism (organizations enhancing each other's fates) to explain evolution (Baum & Singh, 1994: 347). More specifically, density overlap—or domain similarity, in the terminology of Van de Ven (1976)—which refers to overlap in the resource positions of sets of organizations, recognizes that competition exists within networks (Baker, Faulkner, & Fisher, 1998). Baum and Singh elaborated on this idea: "The potential competition experienced by an organization in organizational niche i due to the presence of an organization in organization niche j, is measured by the . . . overlap between organizational niches i and j" (1994: 351). In other words, when organizations occupy the same niche or overlapping niches, there is a high potential for competition for the underlying resources. Consider Baum and Singh's (1994) study of day care centers. The key resource that these businesses compete over is children of various age groups or segments. One class of day care centers may only enroll infants, while another class may only enroll toddlers. The niche of the former day care center is infants, while the niche of the latter is toddlers. There is no overlap between these two niches and hence no competition between the day care centers—the intersection of the organizational capabilities is null. Now consider a third center that enrolls both toddlers and infants. The resulting niche of this business (toddlers and infants) overlaps with the niches of both of the aforementioned businesses and hence competes with both of them.

We consider density overlap as our first explanation because console manufacturers compete for video game developers to release titles for their specific platforms. A platform's niche thus consists of the set of developers and the titles that they have released for the platform. This conceptualization of a niche is akin to that of Podolny and Stuart (1995), who defined technological niche at the level of the individual innovation and included “the identities of the actors associated with the focal innovation and the innovations to which the focal innovation is connected” (1995: 1226). In our conceptualization, games and platforms are the linked innovations. Game developers, however, often do not exclusively release titles to a platform. In fact, a developer’s set of resources may be spread across multiple platforms. Two platform niches may overlap because a developer may release titles for two or more consoles.

For example, consider the case of two platforms and three developers. The niche of platform A is the set of values (1, 0.5, 0.75), and the niche of platform B is the set (0, 0.5, 0.25). Developer 1 releases titles exclusively to platform A; developer 2 releases half its titles to platform A and half to B; and developer 3 releases 75 percent of its titles to platform A. The potential for competition between the two platforms is manifest. Platform A is clearly more crowded (that is, $1 + 0.5 + 0.75= 2.25$) than platform B ($0 + 0.5 + 0.25 =0.75$). When viewed from the perspective of a game developer assessing alternative niches defined by different consoles (different platforms), niche crowding, which exists when many other developers have launched titles for a console, is less favorable. More importantly, a high degree of overlap between two consoles suggests more developers that support both consoles. If one platform begins to attract more developers, the remaining developers can switch, shifting their title releases, further intensifying competition.

On the one hand, density overlap has been shown to influence organizational mortality at both population and segment or niche levels (Baum & Singh, 1994). On the other hand, high levels of niche crowding have been shown to increase competition and the rate of innovation for organizations that choose to remain and compete (Stuart, 1999). Stuart found that such organizations’ rates of research and development investments increase. However, Greve (2000) found strong empirical support for the influence of density overlap on market *entry* choices in Tokyo banking. Our expectation was that density overlap would be a significant explanatory variable for preferential linkage in the context of new network ties: if a developer finds itself with considerable domain similarity to other developers vying for attention on a given platform, the developer will most likely avoid launching another game title in this highly crowded and competitive console niche. The developer will do so because of a sense that its games will not be seen as special or differentiated in such cases by end consumers.

We expected that developers selecting platforms to link with at a given time will more likely select those with lower density overlap. A platform whose position is structurally similar to another platform’s position—in that they draw on overlapping sets of complementors as key resources—will be at a disadvantage for attracting new game titles. Thus, we expected:

Hypothesis 1. Ceteris paribus, a video game developer is less likely to link with a technological platform that has high network overlap density at the time of the developer’s linking decision.

Platform Embeddedness as a Driver of Preferential Linkage

It is generally accepted that organizational actions are embedded in networks of relationships (Granovetter, 1985, Gulati & Gargiulo, 1999; Powell et al., 2004). Detailed case studies and empirical research support this view. Elaborating on the general concept of embeddedness, Uzzi (1996) focused on structural embeddedness. Using the profile of input-output between manufacturers and contractors in vertical relationships to construct his index of second-order coupling, he showed that structural embeddedness impacted mortality in the apparel industry. Designing and launching video games is based on formal and informal agreements between console manufacturers and game developers. The coordination of product launches is a distinguishing characteristic of competition and success in high-tech settings (Arthur, 1989; Shapiro & Varian, 1999). These are not atomistic market-based transactions but rather, reflections of prior ties and patterns of embeddedness (Arthur et al., 1997). Further, as Marsden noted: “Embeddedness refers to the fact that exchanges within a group . . . have an ongoing social structure [that operates] by constraining the set of actions available to the individual actors and by changing the dispositions of those actors toward the actions they may take” (1981: 1210).

In our setting, embeddedness is manifested through the pattern of distribution of titles offered by the different developers for the different platforms. Thus, in the terminology of Uzzi (1996), a platform is tightly coupled when its set of video games is offered by a small number of developers. This situation also implies that the platform is not widely supported by the video game developer community. Uzzi (1997) further noted that high embeddedness could paradoxically reduce an organization’s ability to adapt to new requirements. In an industry characterized by frequent technological changes (Brown & Eisenhardt, 1997), high platform embeddedness may be a liability, as both a console manufacturer and its set of developers may be locked into their current offerings and may not recognize the need for the consoles to adapt to changing requirements. Thus, tightly coupled platforms appear unattractive to new developers, who may seek more central positions with platforms that are loosely coupled. Thus:

Hypothesis 2. Ceteris paribus, a video game developer is less likely to link with a technological platform that has high platform embeddedness at the time of the developer's decision.

Platform Dominance as a Driver of Linkage

We now move from macro network attributes to a focus on characteristics of platforms. Our third explanation is based on the relative market positions of platforms within the network of video game developers and game console producers. Our expectation is that developers will be attracted to the platform that is dominant at a given point in time. There is extensive support for the argument that market dominance creates strong positive feedback, giving rise to increasing returns to scale (Arthur, 1989). As Katz and Shapiro wrote, "Once a certain system is chosen, switching suppliers is costly because new relation-specific investments have to be made. In such a situation, systems that are expected to be popular—and thus have widely available components—will be more popular for that very reason" (1994: 94). Thus, the dominant platform attracts more complementors.

Organization theorists have also supported the notion that companies are attracted to the dominant (or more central) actor in networks for reasons of legitimacy and stability (Oliver, 1990), which reduce market uncertainty. Such an attraction leads to the creation of "hubs" in large networks (Albert & Barabasi, 2002), because new links appear to be driven by the number of previous links to a node. This process gives credence to the concept of the "fit get rich," whereby the fittest (most dominant) node "will inevitably grow to become the biggest hub" (Barabasi, 2002: 103). This concept is consistent with the view that positive feedback occurs in networked settings (Arthur, 1989).

Platform dominance through positive feedback effects in networks is also at the center of the recent U.S. Department of Justice case against Microsoft. Most witnesses for the government argued that Microsoft enjoyed an advantage owing to the strength of complements available for the Windows platform. The testimony at the Microsoft trial by John Soyring, director of network computing of IBM, is typical of the arguments:

> Microsoft. . . has benefited from a . . . cycle that also tends to be self-perpetuating in the absence of some industry advance that undermines it. The large installed base of Windows has encouraged ISVs [independent software vendors] to develop a large number of applications for Windows, which has led to increased demand for Windows. This, in turn, has further increased the incentive for ISVs to develop applications for Windows. In fact, users, ISVs, and PC suppliers—recognizing the opportunity offered by Windows' large installed base—know that the most popular applications will be written for Windows. In fact, given the relative size of the Windows installed base, no PC application can achieve wide distribution—that is, be a "best seller"—that is not offered on Windows.[1]

The preliminary legal "findings of fact" upheld the view that Microsoft enjoyed high barriers to entry (McKenzie, 2000) because a given independent software developer would be disinclined to write applications for any platform other than Windows.

Thus, the dominant platform at any given time will preferentially attract new links with existing and new game developers. Developers are attracted to the dominant platform because it offers the biggest potential market for their new games. Moreover, by choosing to link with the dominant platform, the developers also minimize uncertainty about the future viability of other platforms, especially under conditions in which the dominant platform may enjoy increasing returns and positive feedback effects (Shapiro & Varian, 1998). Thus:

Hypothesis 3. Ceteris paribus, a video game developer is more likely to link with a technological platform that has a dominant market position.

Platform Age as a Driver of Linkage

Our fourth assertion, which concerns platform age, is a counterbalance to the platform dominance argument made above. One way in which dominant platforms with extensive network connections lose their attractiveness is that their functionality declines relative to that of newer platforms. In high-velocity environments, rapid recognition and adaptation to technological shifts is a key requirement (Brown & Eisenhardt, 1997). When new platforms are introduced to take advantage of significant improvements in computer processing power, speed, and memory, every developer faces a new choice: to launch products on this new platform, or not do

[1] http://www.usdoj.gov/atr/cases/f2000/2054.htm; accessed July 2003.

so. Platform innovations also challenge the entrenched dominant position. If a platform's dominance is perpetual and its supporting positive feedback cycle cannot be broken, strategic innovations and maneuvering have no role. But leadership positions in high-velocity environments change rapidly, giving rise to serial monopoly, a situation of "winner-take-all, but only for a while" (Liebowitz & Margolis, 2001: 137). Both console makers and game developers need to adjust to new innovations, and patterns of interorganizational relationship are not static.

Platform evolution through technological development shapes network structure and the path of its evolution. A summary of Richard Schmalansee's January 1999 testimony at the Microsoft trial contains these observations:

> History shows that application software emerges quickly for promising software platforms. In the early 1980s, for example, independent software vendors wrote applications for DOS despite the fact that many consumers were using CP/M, and they wrote for Windows in the early 1990s even though most consumers were using DOS. Today software developers are writing applications for Linux and the PalmPC platform, and developers have shown renewed interest in the Apple Macintosh now that Apple has released a new range of computers that is proving popular with consumers.[2]

Technological innovations in video games have introduced new platforms, new competitors, and new competitive moves (Gallagher & Park, 2002). Every new platform introduction potentially alters the established network of links between developers and console manufactures. A new design, the result of significant improvements in the basic functionality of consoles, can destroy the competencies of some developers while enhancing the competencies of others. The introduction of a new platform thus poses a new challenge for all developers: should they continue to offer new titles for existing platforms or alter their positions to take the new platform into consideration?

Video game developers are acutely aware that their network positions rely on staying at the cutting edge of technology and that they will be considered unattractive if they fall behind. This logic is particularly important in a networked market in which the platform leader may simply reallocate marketing resources to those developers that have stayed up-to-date with technology developments. The introduction of a new video game platform also offers new and exciting ways for game developers to showcase designs that capitalize on the new platform's faster speed, superior graphics, and enhanced multimedia functionality. The developers also may be drawn to ride the marketing blitz involved in launching a new platform with tie-in promotions with lead game titles. Thus:

> *Hypothesis 4. Ceteris paribus, a video game developer is more likely to link with a newer technological platform.*

METHODS

Research Setting: The Programmable Video Game Network

As noted above, the setting we chose for our examination of patterns of linkage formation in networks is the complex and dynamic video game sector. It is a nexus of U.S. organizations that offer video game consoles and video games (Schilling, 2003; Shankar & Bayus, 2003). Each launch of a new game for a given console represents a network link. The console choices that developers make over time differentiate network positions and influence customer preferences. Programmable home video game consoles (hardware) and games, or game titles (software), are complements, since one is of little value without the other (Brandenberger & Nalebuff, 1996). The dynamism of this setting comes from the frequency with which manufacturers launch new generations of consoles with enhanced capabilities (for instance, faster processors, better graphics, and more memory) that offer game developers additional capabilities to exploit (Gallagher & Park, 2002). Thus, this setting provides a rich context in which to examine questions of preferential linkage driving network evolution.

Data

We assembled a primary data set of unique developer-title-console combinations in monthly intervals to capture our network. We first identified 671 developers and 8 platforms. Then, our intent was to consider every title released by every developer for every console during the time frame of our study, in the spirit of Salancik's (1995) call for attention to the "organization of individual actions." This level of detail was required because every video game title represents new content. Consider the Madden NFL series of titles developed by Electronic Arts. A new Madden NFL title is released every year but with different football players and updated athlete statistics. Similarly, Final Fantasy, developed by Square Enix, was first developed and released in 1990. Since then, many dif-

[2] www.microsoft.com/presspass/trial/jann99/01–11schmal.asp. Accessed July 7, 2003.

ferent Final Fantasy titles have been released with new storylines and new game challenges.

Our data captured the sequence of developer-platform links month by month. For example, Electronic Arts launched "Need for Speed: Hot Pursuit 2" for the GameCube console in September 2002, for the Xbox console in October 2002, and for the PlayStation2 console in November 2002. We thus treated each launch as a distinct link since a game for one type of console cannot be used on a different type, just as Microsoft Word for Windows cannot be used on a Macintosh. At the end, our final sample consisted of 2,815 releases, each representing a unique combination of game developer, console, and title, spanning the eight years from January 1995 through December 2002, or 96 months. In the Appendix, we describe in detail the steps we took to assemble and validate the dataset.

TABLE 1
Release Sequence for Two Game Titles

Game Title	Release Date	Choice Set	Console Choice
"Madden NFL 2002"	8/2001	GameCube PSOne PlayStation2 Xbox	PlayStation2
"Madden NFL 2002"	10/2001	GameCube PSOne Xbox	Xbox
"Madden NFL 2002"	11/2001	GameCube PSOne	GameCube
"Need for Speed: Hot Pursuit 2"	9/2002	GameCube PSOne PlayStation2 Xbox	GameCube

Measures

Linkage. We defined the dependent variable, a developer's choice to release a game title for a given console, as categorical and time-varying. Specifically, $y_{ij,t}$ equaled 1 if developer *i* on date *t* released a game title for console *j,* where *j* was a console in the developer's choice set, J_{it}. For the remaining consoles in the developer's choice set, $y_{ij,t}$ equaled 0. Table 1, which traces the dates on which Electronic Arts released "Madden NFL 2002," illustrates the coding of the dependent variable. As the table shows, the set of console choices $(J_{i,t})$ for Electronic Arts' first release of "Madden NFL 2002" included GameCube, PSOne, PlayStation2, and Xbox. The game was first released by Electronic Arts *(i)* for PlayStation2 *(j)* in August 2001 *(t)*; hence, $y_{ij,t}$ is 1. For the three remaining consoles, $y_{ij,t}$ is 0. For the second release of "Madden NFL 2002," the choice set excluded PlayStation2 because developers never release the same title to a console twice. "Madden NFL 2002" was released for Xbox in October 2001. For the final release of this game, its release for GameCube in November 2001, the choice set excluded both PlayStation2 and Xbox. As this example demonstrates, the choice set was not constant. Table 1 further illustrates that for a new and previously unreleased title ("Need for Speed: Hot Pursuit 2"), the choice set was reset and would subsequently decrease as the same title was released for additional consoles.

Table 1 illustrates one way in which the choice set varied across time as a game title was sequentially released on multiple platforms. However, as a console became obsolete, it was permanently removed from all choice sets and calculations. We determined from our sample the *last* date t^* that *any* developer released a new title for console *j.* It is obsolete if the release date *t* of a different title is greater than t^*. For example, developers in 2000 stopped releasing titles for the Sega DreamCast console. For this reason, DreamCast was not a console in any of the choice sets for "Madden NFL 2002" and "Need for Speed: Hot Pursuit 2" in 2001. Consequently, we removed all the titles previously released for the obsolete console from the network; such removal was also in line with Burt's (2000) suggestion that a network evolves via the decay of the actors (here, console manufacturers) and links (game titles). Finally, the introduction of a new console increases the number of consoles in a developer's choice set.

We also recognized the exit of game developers by including in all calculations only active developers and their titles. We removed a developer and its titles from the network after it stopped releasing new titles. From our sample, we were able to determine the *last* date on which a developer released a new title for any console. Finally, for all calculations, we measured all model predictors just prior to date *t.*[3]

Platform and developer age. We defined platform age in months on the release date *t* of a game title as the difference between that date and console *j*'s launch date. We also defined developer age in months as the difference between date *t* and the date of developer *i*'s first title release for any console.

Platform and developer dominance. We defined platform *j*'s dominance on date *t* as the number of titles released by all developers for the platform divided by the total number of titles in the network. Similarly, we defined developer *i*'s dominance on date *t* as the number of titles released by

[3] In this text, we omit the subscript *t* to enhance expositional clarity.

the developer for all platforms divided by the total number of titles in the network.

Platform overlap density. We used overlap density (Baum & Singh, 1994) to measure the potential competition faced by developers in the platform j niche. The overlap density for platform j on date t was given by

$$Overlap\ Density_{j,\,t} = \sum_i p_{ij,\,t} + \sum_{k \neq j} w_{kj,\,t}\left(\sum_i p_{ik,\,t}\right).$$

The term $p_{ij,\,t}$ denoted the proportion of titles that developer i released to platform j. Summing over all developers thus denoted the total number of developers releasing titles for platform j. Similarly, $p_{ik,t}$ denoted the proportion of titles that developer i had released to platform k. Summing over all developers thus distinguished the total number of developers releasing titles for platform k from the total releasing titles for platform j.

Baum and Singh (1994) defined $w_{kj,\,t}$ as the degree of overlap between two market niches. Market niches, in our research setting, corresponded to two platforms, j and k. The platform overlap degree or weight was defined on the interval $0 \leq w_{kj,\,t} \leq 1$ and denoted the extent to which two platforms attracted the same developers and thus overlapped. If $w_{kj,\,t}$ equaled 0, two disjoint sets of developers had released titles for platforms k and j, and the intersection of the two sets was null. If $w_{kj,\,t}$ was 1, developers were releasing titles for both platforms. Adding the product of the overlap weight and the total number of developers that released titles for k to the total number of developers that released titles for j thus denoted the total competition confronting a developer releasing a title to platform j.

Many specifications of overlap weight exist, with the Euclidean metric being the most common (e.g., Burt & Carlton, 1989). However, as Baum and Singh stated, "The overlap weight [is], in general, asymmetric" (1994: 351–352)—or, $w_{jk,\,t}$ is usually not equal to $w_{kj,\,t}$. The Euclidean metric is symmetric. We thus used Sohn's (2001) overlap metric, defined as

$$w_{kj,\,t} = \frac{\sum_i a_{ik,\,t}\ min(a_{ik,\,t}, a_{ij,\,t})}{\sum_i (a_{ik,\,t})^2},$$

where $a_{ij,\,t}$ and $a_{ik,\,t}$ are the numbers of titles released by developer i to platforms j and k, respectively.

Developer prior ties. We defined prior ties (Granovetter, 1985) on date t as the number of titles that had been released by developer i for platform j by that date.

Platform embeddedness. We defined platform j's embeddedness using Uzzi's (1996) second-order coupling. Let $P_{ji,\,t}$ denote the proportion of platform j's video games released by developer i on date t for platform j. Platform embeddedness was calculated as $\sum_i (P_{ji,\,t})^2$ and varied between 0 and 1. A value close to 1.0 for this index meant that a platform was tightly embedded with (supported by) a few developers. A value close to 0 meant a platform was loosely embedded with (supported by) many developers.

Developer embeddedness. Similarly, we defined developer embeddedness using Uzzi's (1996) first-order coupling. Let P_{ij} denote the proportion of the developer i's titles released for platform j on date t. Developer embeddedness was $\sum_j (P_{ij,\,t})^2$. A value of 1 meant that 100 percent of a developer's titles were being released for a single platform, and a value close to 0 meant a developer was releasing titles for many platforms. Table 2a summarizes means, standard deviations, and zero-order correlations for the variables.

Modeling Console Choice and Network Evolution

Our unit of analysis was a developer's choice to release a title for a specific console. We thus used McFadden's (1974) multiprobability discrete choice model, a generalization of "logit" and multinomial logit models, to test our hypotheses. A multiprobability model is commonly used when the dependent variable consists of a choice from an unordered set of alternatives. For example, Greve (2000) used the McFadden model to estimate market entry choices in Tokyo that were based on attributes of alternative locations within a consideration set. Greene and Hensher (1997) used the multiprobability model to analyze choice of mode of transportation from a set of four alternatives. Powell, White, Koput, and Owen-Smith (2004) also used the McFadden model, to model partner choice in the biotechnology sector.

All covariates here time-varying because they represented the network structure on any date. Moreover, a particular characteristic of McFadden's multiprobability discrete choice model is that it requires as input information on a choice and all the alternatives available at the time of the choice. Greene and Hensher's (2003) analysis of 210 people's choice of travel mode illustrates the model. Their data set contained 840 observations because each person chose from a fixed choice set (air, train, bus, car). Hence, each of the 210 choices was associated with three "nonchoice" observations. Following Greene (2003: 730), we present both sets of descriptive data—one for the choices (the study sample) and another for the choices and all alternatives (the estimation sample).

TABLE 2
Summary Statistics and Correlations

(2a) Study Sample

Variable[a]	Mean	s.d.	1	2	3	4	5	6	7
1. Platform age	33.61	21.14							
2. Platform dominance	0.36	0.28	.55						
3. Platform overlap density	118.11	33.43	.32	−.19					
4. Platform embeddedness	0.04	0.05	−.34	−.21	−.32				
5. Developer age	29.03	25.27	−.01[b]	−.34	.37	−.05			
6. Developer dominance	0.01	0.02	−.15	.12	−.31	.39	.18		
7. Developer embeddedness	0.70	0.26	.17	.44	−.28	.02[b]	−.61	−.09	
8. Developer prior ties	4.45	6.23	.37	.25	.13	−.13	.33	.40	−.12

(2b) Estimation Sample[c]

Variable	Mean	s.d.	1	2	3	4	5	6	7
1. Platform age	38.91	26.81							
2. Platform dominance	0.22	0.23	.43						
3. Platform overlap density	113.97	33.24	.16	−.04					
4. Platform embeddedness	0.07	0.16	−.35	−.26	−.14				
5. Developer age	31.83	25.60	.16	−.13	.26	−.03			
6. Developer dominance	0.01	0.02	−.08	.08	−.22	.09	.25		
7. Developer embeddedness	0.68	0.26	−.14	.12	−.20	.02	−.61	−.14	
8. Developer prior ties	2.28	5.46	.27	.36	.08	−.12	.31	.35	−.20

[a] n = 2,815. Correlations are significant at $p < .05$ or less owing to the large sample size, except as noted.
[b] $p < .10$.
[c] n = 11,931. All correlations are significant at $p < .05$ or less owing to the large sample size.

Our study sample contained 2,815 unique developer-platform-title observations. Our estimation sample included, in addition to the study sample, all the consoles not chosen in each of the 2,815 cases. However, the choice set for a developer's title was not fixed. Consider again the sequential launches of "Madden NFL 2002" (Table 1). For the first launch of the game, the choice set consisted of four consoles; this release thus generated four observations (that is, one for PlayStation2 and three for the other consoles). For the second launch of "Madden NFL 2002," three console choices generated three observations (one for Xbox and one each for PSOne and GameCube). Moreover, the maximum number of consoles in a choice set varied over time as obsolete consoles were removed from the choice set. Consequently, as noted above, each of the 2,815 title releases created one or more additional observations that consisted of all active consoles for which the developer had not released the title. The final estimation sample thus consisted of 11,931 observations. Summary statistics for the study and estimation sample are found in Tables 2a and 2b.

We should also note the specification of model covariates in McFadden's multiprobability discrete choice model: covariates are classified as varying with the choices in the choice set (here, platform attributes) and as not so varying (here, developer attributes). Again, the August 2001 release of "Madden NFL 2002" in Table 1 is illustrative. Each of the four consoles in the choice set seen there was launched by the manufacturer on a different date, so platform age, for example, differed for each console. The consoles, however, represented choices that confronted Electronic Arts on August 2001. Consequently, developer age, for example, did not vary across the choices. As Greene stated, "Terms that do not vary across alternatives – that is those specific to the individual – fall out of the probability. One method [is to create] . . . a complete set of interaction terms" (2003: 720) (between the choice and individual covariates). In other words, to incorporate both platform and developer attributes simultaneously, developer attributes needed to be specified as "interaction" terms in the multiprobability model. Although the inclusion of many interaction terms complicated the time-varying model, omitting them in the model specification would have implied that all developers were homogeneous. The coefficients for the "main effects" were robust only given inclusion of these interactions, which is akin to the inclusion of control variables in multivariate regressions.

Hence, the probability that developer i would release a title for platform j instead of for the other platforms in $J_{i,t}$ was a function of a vector of platform attributes $(X_{j,t})$ that varied among all platforms in $J_{i,t}$; developer-specific dyadic attributes $(W_{ij,t})$; and a vector of developer attributes $(Z_{i,t})$ that did not vary. The estimation model was:

$$Probability\left(y_{ij,t} = 1 \mid \sum_{j_{i,t}} = 1, X_{j,t}, Z_{i,t}, W_{ij,t}\right) = \frac{exp(\beta X_{j,t} + \gamma Z_{i,t}X_{j,t} + \lambda W_{ij,t}X_{j,t})}{\sum_{j_{i,t}} exp(\beta X_{j,t} + \gamma Z_{i,t}X_{j,t} + \lambda W_{ij,t}X_{j,t})}.$$

We estimated two models. Model 1 (homogeneous developers) includes only platform attributes: overlap density, embeddedness, platform dominance and platform newness. In model 2, we relaxed the assumption of homogeneity by incorporating the four developer characteristics.

Both models were estimated with Stata 8.2. The standard errors were based on the coefficients from the estimation. However, we also estimated and present odds ratios for model 2, obtained by taking the exponential of the coefficients. An odds ratio greater than 1.0 denoted an increase in probability for higher values of the coefficient, and a value lower than 1.0 denoted a decrease.

RESULTS

Table 3 presents estimations for the two models, which are both significant (pseudo-R^2 = .24 and .52). The between-model change in chi-square is also significant. The high pseudo-R^2 in model 2 adds credence to the results, despite the large sample size. We used the full model, with both platform and developer attributes, for testing the four hypotheses, assessing support for a hypothesis through the main effects (rows 1–4).

Our first hypothesis, predicting that density overlap will have a negative effect on network linkage, was supported ($b_1 = -0.02$, odds ratio = 0.98, $p < .05$). Developers were less likely to launch products in a crowded space in which other devel-

TABLE 3
Results of Multiprobability Regression Analysis

		Model 1		Model 2		
Variable	Hypothesis	*b*	s.e.	*b*	s.e.	Odds Ratio
Platform overlap density	1	0.01***	0.00	−0.02*	0.01	0.98
Platform embeddedness	2	−6.74***	0.76	−39.58***	6.71	0.00
Platform dominance	3	3.78***	0.18	2.83*	1.14	16.86
Platform age	4	−0.04	0.00	−0.05***	0.01	0.95
Platform overlap density × developer age				0.00***	0.00	1.00
Platform overlap density × developer dominance				0.41	0.35	1.50
Platform overlap density × developer embeddedness				0.02†	0.01	1.02
Platform overlap density × developer prior ties				0.00†	0.00	1.00
Platform embeddedness × developer age				0.23**	0.07	1.26
Platform embeddedness × developer dominance				−208.46***	42.30	0.00
Platform embeddedness × developer embeddedness				19.34**	7.10	2.5e8
Platform embeddedness × developer prior ties				24.75***	1.19	5.6e10
Platform age × developer age				−0.00**	0.00	1.00
Platform age × developer dominance				−3.95***	0.37	0.02
Platform age × developer embeddedness				0.03*	0.01	1.03
Platform age × developer prior ties				−0.00**	0.00	1.00
Platform dominance × developer age				0.08***	0.02	1.09
Platform dominance × developer dominance				−203.18***	27.24	0.00
Platform dominance × developer embeddedness				−0.15	1.20	0.86
Platform dominance × developer prior ties				0.47***	0.10	1.61
Log-likelihood		−2,954.95		−1,878.85		
Likelihood ratio chi-square		1,898.48		4,050.68		
Pseudo R^2		.24		.52		

† $p < .10$
* $p < .05$
** $p < .01$
*** $p < .001$

opers and titles had strong presences. The interpretation of the odds ratio is that for a unit increase in density overlap, the likelihood of attachment decreases by 2 percent (1.00 − 0.98). We then tested Hypothesis 2, on platform embeddedness through coupling. The effect was strong and in the expected direction ($b_2 = -39.58$, odds ratio: 0.00, $p < .001$), indicating that developers were unlikely to release titles to a platform that was already tightly connected to few developers.

Hypothesis 3 focuses on platform dominance. As a predictor, platform dominance was strong and significant ($b_3 = 2.83$, odds ratio = 16.86, $p < .05$). It is noteworthy that the impact of platform dominance on developers' platform choice was much stronger than the impact of competition within a platform seen through overlap density. The implication is that for a 10 percent increase in platform dominance, the odds of attachment increases 1.686 times (16.86/10, given the range of platform dominance from 0.0–1.0). Our Hypothesis 4, stating that in a fast-changing technology setting, developers will be drawn to launch new titles for the most innovative gaming platform, received strong empirical support ($b_4 = -0.05$, odds ratio = 0.95, $p < .001$). A developer was more likely to release a title for a newer platform, because the odds of releasing a title for a platform decreased by 5 percent (1.00 – 0.95) for a unit increase in the difference between the title's launch date and the platform's launch date. For the four hypotheses collectively, empirical support was strong.

Further Validation

The analyses reported thus far measured platform dominance as the number of new game launches instead of as the more refined measure, share through sales in units. To further validate our results, given the importance of platform dominance, we obtained additional data from the NPD Group®, an authoritative industry source whose data has been previously used by researchers (Shankar & Bayus, 2003). This commercial database, spanning the years 1995–2002, contains the same information structure (that is, developer, title, console, launch date) as our event sample. The NPD database also has yearly data on the total copies of a game title–platform sold (that is, units shipped) to retailers (for instance, Best Buy®) and the unit share of a title for a specific console. These additional data are important because, although all new titles are priced similarly at the time of launch, popular titles linked to a specific platform may have more influence over developers' choices to release to a console.

Unfortunately, this database contains observations for only the top 100 titles. Nonetheless, we were able to determine by performing a few calculations that the top 100 titles accounted for approximately 68 percent of industry units shipped. We then recalculated platform dominance using units shipped instead of new title counts to test the robustness of our results. We reestimated the model with yearly data by recalibrating the measures at yearly resolutions.

The results, presented in Table 4, were remarkably consistent with those for the original model: the pseudo-R^2 was .50 for the validation (NPD) model versus .52 for the original model. This consistency is not surprising, given that the average correlation over the eight study years between platform dominance (calculated using the number of new titles, our primary data), and unit sales (calculated using the number of shipped titles, from NPD data), is .85. The implication is that the sales level of platforms is significantly influenced by new title launches, the building block for network evolution that we studied here. This correlation also provides further confidence in the quality of our data and the veracity of our results.

Specifically, the coefficient for platform dominance was significant at a better level of statistical precision ($p < .05$, original analysis; $p < .01$, validation analysis). Three of the four platform attribute effects were significant; density overlap's significance as a main effect was not supported in the validation analysis. The developer attribute "controls" were broadly supported, and the signs of the developer attribute coefficients (with the exception of the interaction of platform age with developer prior ties) were identical—albeit the Z-values obtained using the NPD data were smaller than those from our primary data.

The degree of density overlap (and of other effects) observed via monthly "snapshots" of network evolution might be greater than what would be observed if our data were aggregated yearly. Since the validation analysis only included developers with at least one title in the top 100, the level of competition was underrepresented. It could also be that the top 100 titles had different impacts in the limited analysis than they did when all the titles were used. Nevertheless, we believe that the supplementary analyses enhanced the robustness and validity of our results.

From Statistical Estimation to Network Visualization

We sought to shed further light on network evolution by employing complementary network modeling and visualization approaches (Albert & Bara-

TABLE 4
Results of Validation Multiprobability Regression Analysis

Variable	Hypothesis	Model 1		Model 2	
		b	s.e.	*b*	s.e.
Platform overlap density	1	0.02***	0.00	0.01	0.01
Platform embeddedness	2	−19.96***	2.46	−52.52***	13.44
Platform dominance	4	4.80***	0.21	3.83**	1.12
Platform age	4	−0.12***	0.02	−0.28**	0.09
Platform overlap density × developer age				0.00	0.00
Platform overlap density × developer dominance				0.42	0.27
Platform overlap density × developer embeddedness				0.00	0.01
Platform overlap density × developer prior ties				0.00***	0.00
Platform embeddedness × developer age				−0.07	2.17
Platform embeddedness × developer dominance				−73.08	260.89
Platform embeddedness × developer embeddedness				34.05*	13.31
Platform embeddedness × developer prior ties				13.38***	1.13
Platform dominance × developer age				0.55**	0.16
Platform dominance × developer dominance				−181.76***	48.81
Platform dominance × developer embeddedness				−1.10	1.17
Platform dominance × developer prior ties				0.08	0.13
Platform age × developer age				−0.02	0.01
Platform age × developer dominance				−52.43***	4.70
Platform age × developer embeddedness				0.36***	0.10
Platform age × developer prior ties				−0.00	0.01
Log-likelihood		−2,970.32		−2,109.42	
Likelihood ratio chi-square		2,540.21		4,262.01	
Pseudo-R^2		.30		.50	

* $p < .05$
** $p < .01$
*** $p < .001$

basi, 2002; Powell et al., 2004), using Pajek[4] (Batagelj & Mrvar, 2003) to triangulate statistical estimates of links and network topology structure changes. We mapped linkage in networks by connecting game titles with platforms to derive the topology of the networks over time. Space constraints prevent us from including yearly evolution here, but we include three snapshots. Figure 1 depicts the network at the beginning of our study period, December 1995, which captures the impact of the introduction of Sony PlayStation. Figure 2, for December 2000, captures the impact of the introduction of PlayStation2. Figure 3, for the end of our study period, December 2002, highlights the competitive interaction between Sony and Microsoft with the introduction of the Xbox. In these figures, each circle denotes a unique developer, and each line denotes one or more titles launched by a developer to support a game platform, which is denoted as a square. The sizes of the squares and circles are proportional to the number of titles released by developers (for circles) or written for a platform (for squares).

The three figures taken together allow us to make a few useful points. First, they reinforce the central role of preferential working relationships with few console makers: many developers who supported the original PlayStation also supported the subsequent platform, PlayStation2, introduced by the same organization, Sony. Second, the sequence of graphs demonstrates the strength of platform dominance as an attractor of new titles, with Sony PlayStation shown as receiving many new titles in Figure 1 and Sony PlayStation2 as receiving many in Figures 2 and 3. Third, domain overlap is a central issue; the graphs show many smaller players supporting one platform while the more prominent developers (those with bigger circles) are clearly supporting multiple platforms (Figures 2 and 3). Furthermore, the figures taken together highlight a power shift in the network from platform producers to a few major game developers. By December 2002, more developers were supporting multiple platforms. These Pajek figures reveal useful insights

[4] Pajek is a program for the analysis and visualization of large networks involving many hundreds of vertices. In Slovenian, it means "spider." For a detailed illustration of the ability of Pajek to visualize interorganizational networks, see Powell et al. (2004).

FIGURE 1
Network Topology, December 1995[a]

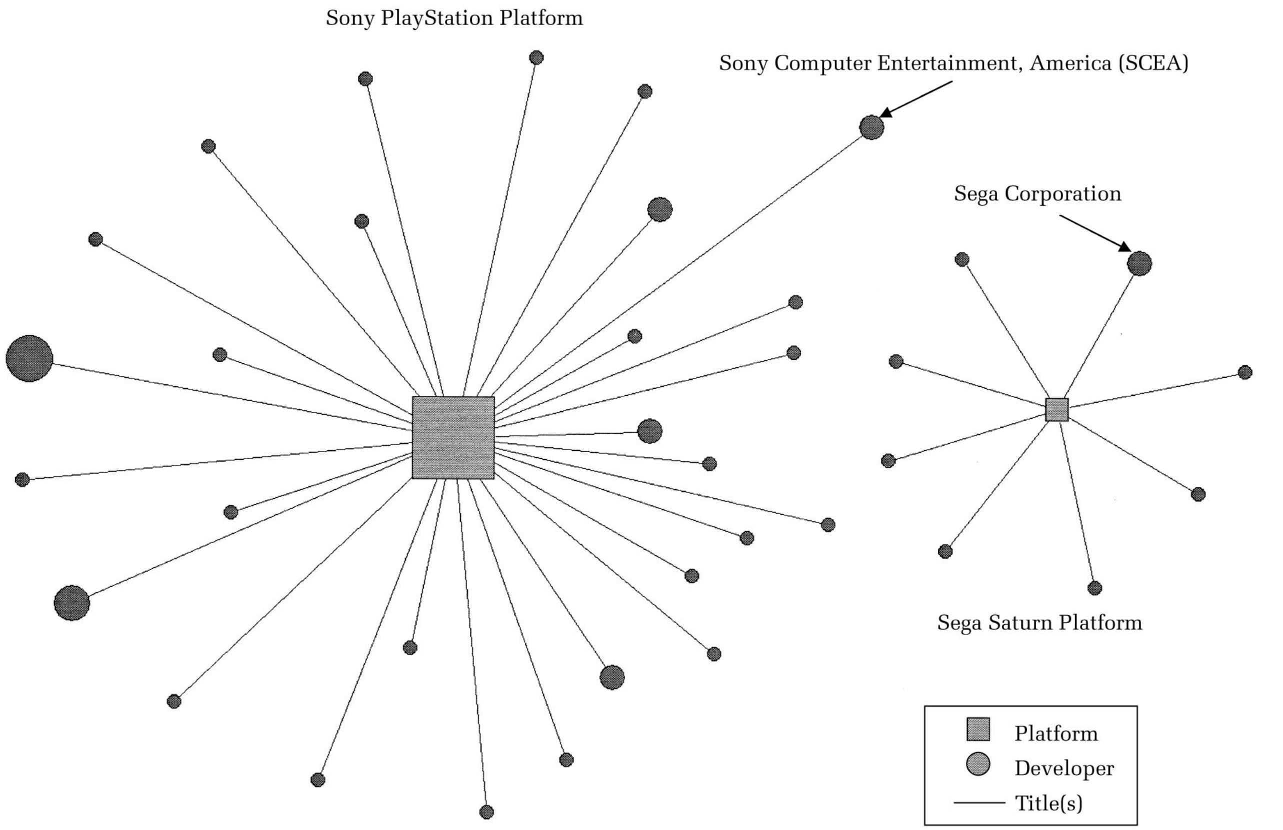

[a] These are summary observations for two nonoverlapping (disjoint) networks. The sizes of the squares and circles shown here are proportional to network activity. For example, the larger a circle, the greater the number of titles released by the developer denoted by that circle.

into the structure and dynamics of networks and complement the statistical analyses.

DISCUSSION

We focused on the dynamics of network links between game consoles and video game developers—two clearly demarcated network roles—over an eight-year period in the U.S. video game sector. We theorized four related explanations for game developers' preferential linkage to specific platforms and found strong empirical support. Results suggest that one cannot understand such networks by focusing on any one role since each developer's success depends on its coordinating choices with those of platform manufacturers (Shapiro & Varian, 1998). The network dynamics of coordination between platform manufacturers and game developers are central since vertical integration is ineffective in such system-based settings. Although some leading console makers like Sony have some of their own titles (see Figure 1), success in such networks is critically based on relationships with complementors. Similarly, game developers need to make critical choices about linking with one or more platforms and adapting these links over time. We also found further corroborating insights from network visualization models.

Specifically, our results show that video game networks evolve, with the formation of links reflecting macro network characteristics (density overlap and embeddedness) and platform characteristics (dominance and newness) when developer characteristics (age, share, coupling, and prior ties) are used as controls to specify a more complete model. Structural embeddedness, seen through the coupling of consoles and video games, emerged as the strongest predictor of linkage. This result adds further credence to Uzzi's (1997) notions of the paradox of embeddedness. Interestingly, our findings are consistent with the literature on network effects in technology systems (Arthur, 1989; Bresnahan & Greenstein, 1999). Firms designing complementary products are drawn to platforms that have broad-based support. Future work may usefully integrate ideas of embeddedness with net-

FIGURE 2
Network Topology, December 2000[a]

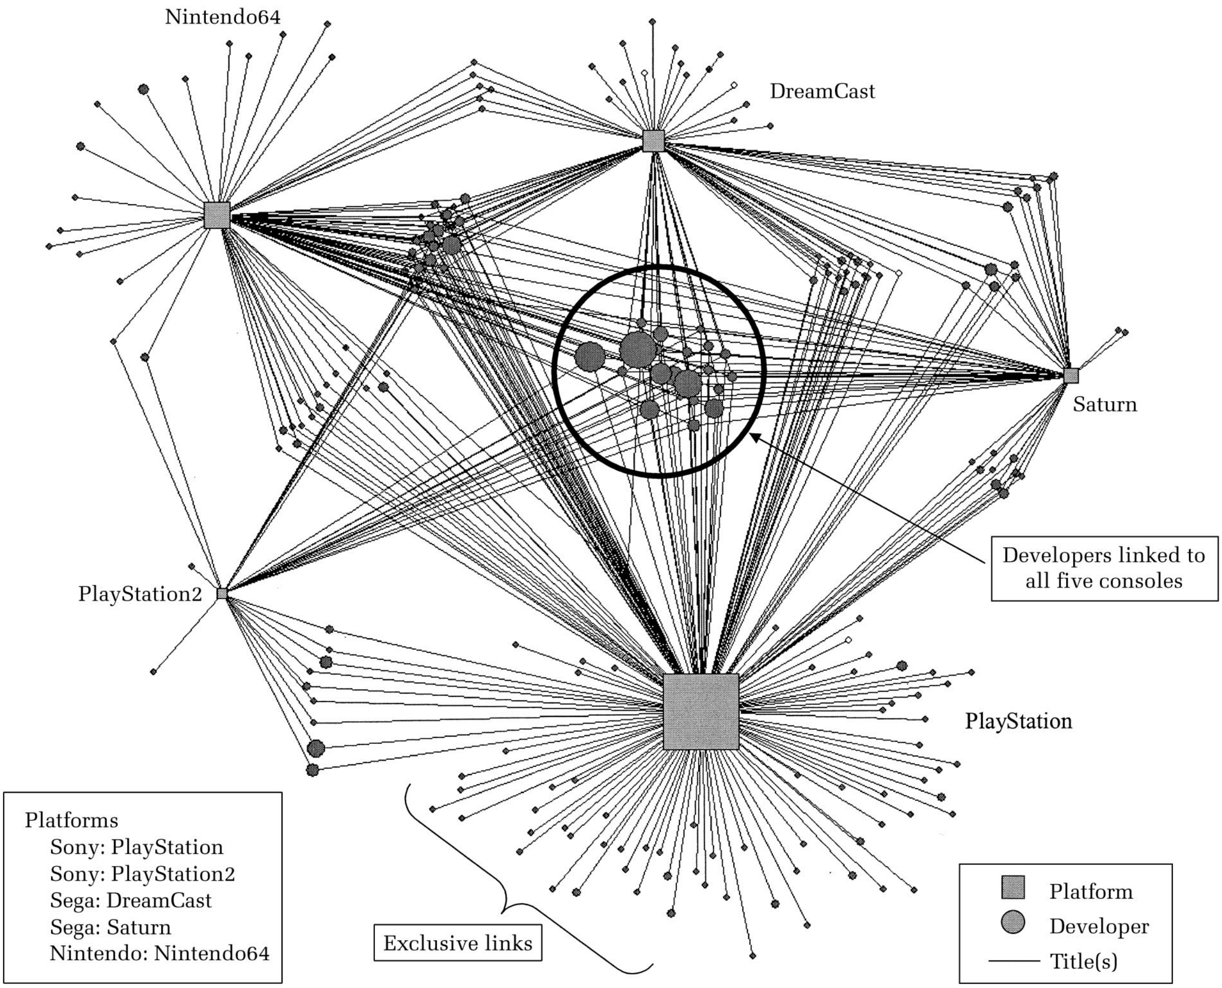

[a] The sizes of the squares and circles shown here are proportional to network activity. For example, the larger a circle, the greater the number of titles released by the developer denoted by that circle.

work effects to develop richer conceptualizations and predictions for later empirical tests.

Platform newness emerged as a significant predictor of attracting new links, lending further credence to the importance of innovation in fast-changing markets (Brown & Eisenhardt, 1997). Developers are drawn to newer platforms, all other things being equal. This is the first large-sample study that we are aware of that focused on the dynamic process of network evolution in a setting in which the technological characteristics changed in significant ways to influence new platform launches. They subsequently impact preferential links among organizations.

In addition to density overlap, extant coupling between platforms and developers, and platform newness, we found that platform dominance influences network links. Physical and technological networks evolve in ways that support the linking logic of "the fit get richer" (Albert & Barabasi, 2002). However, dominance is just one explanation when network evolution over eight years is examined as monthly addition and deletion of nodes and links. Since platform dominance is not permanent (Evans et al., 2001) but dynamic, astute developers strive to balance their commitment of complementary products over an array of platforms with different characteristics and positions in networks.

From a strategic management point of view, our results reinforce the need to understand competition-cooperation in networks. Some interorganizational networks foster innovation through direct and indirect ties (Ahuja, 2000), and other networks may have "small world" characteristics (Davis, Yoo, & Baker, 2003; Kogut & Walker, 2001; Watts, 2003). We treated games as complementary re-

FIGURE 3
Network Topology, December 2002[a]

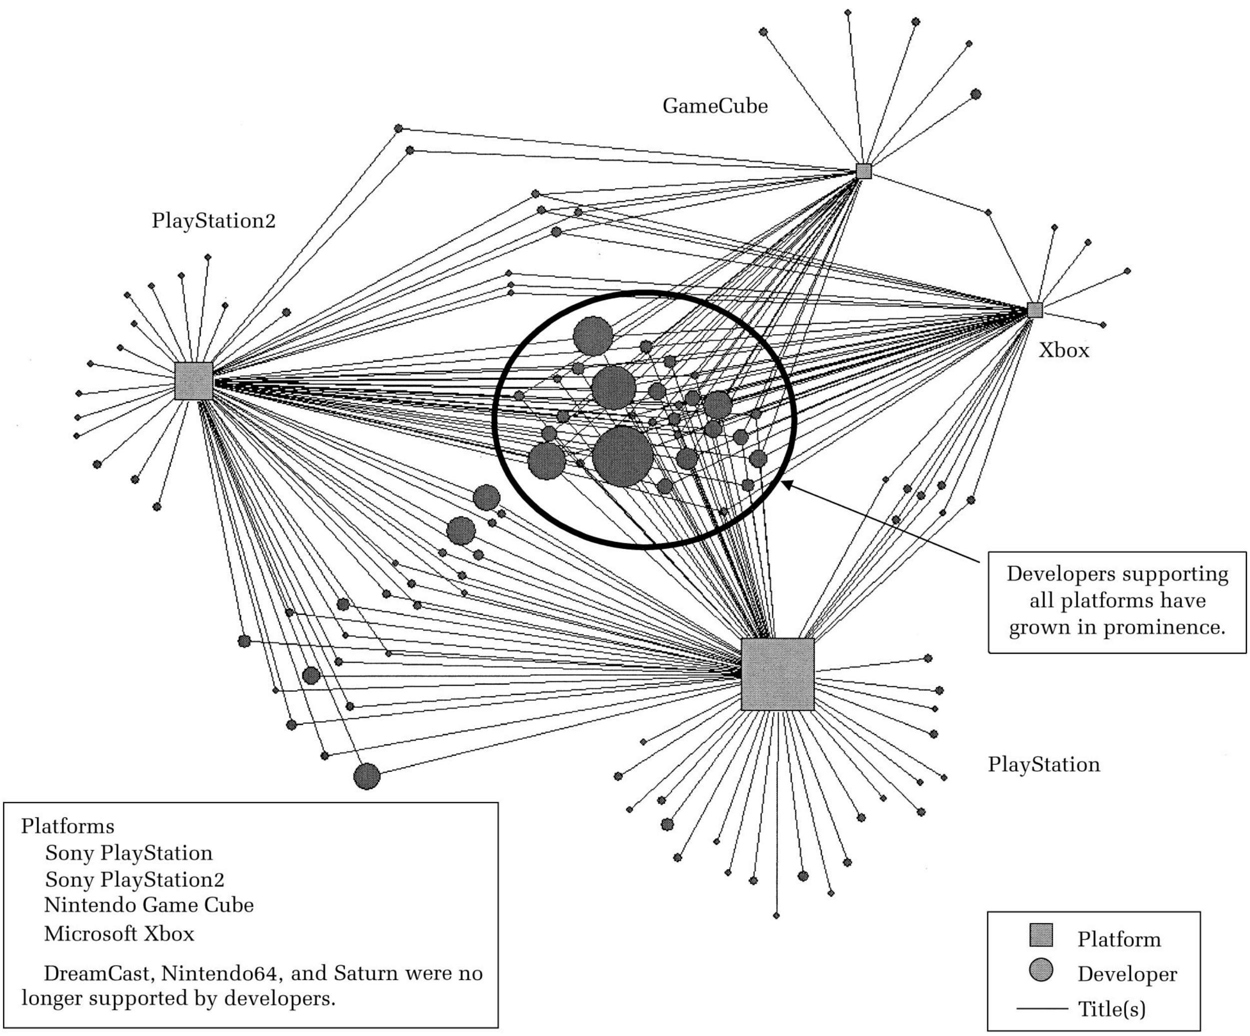

[a] The sizes of the squares and circles shown here are proportional to network activity. For example, the larger a circle, the greater the number of titles released by the developer denoted by that circle.

sources for a platform in the spirit of Podolny's (2001) view of ties as pipes (resource conduits) and ties as prisms (reflecting cues on node quality). However, in other settings, a tie may be only a pipe or only a prism. Furthermore, not all types of interorganizational ties create conduits for resource flows, because competitive interactions find different gateways to resources. In social networks, for example, each actor may be "actively trying to coopt those with whom it has dependencies" (Borgatti & Foster, 2003: 1004). Although bilateral relationships facilitate resource flows (Gulati & Gargiulo, 1999), it is naïve to assume that extended ties through different organizations always serve as conduits for resource flows. We hope that our empirical findings in one type of network reflecting system-based competition stimulate future theorizing about the particular competitive-cooperative nature of these ties and add to the broader stream of research focused on how network characteristics impact network performance (Baum et al., 2003; Powell et al., 2004; Uzzi & Spiro, 2003).

Methodologically, our study makes modest contributions. We went beyond describing network structure to analyze network evolution from a decision-making perspective. Greve (2000) demonstrated the value of models of choice in settings in which decision makers explicitly consider available alternatives in terms of market characteristics. In this study, we used a set of platform and developer attributes to shed light on linkages in networks. Such models explicate the logic of network dynamics as aggregates of individual actions (Salancik, 1995). We also combined statistical estimations of preferential linkage with network visualization approaches to gain insights on network evolution. We believe that such an integrative approach, in line with that of Powell and his colleagues (2004), may be useful for understanding the intricacies of networks. This triangulation allowed us to see patterns that would have remained abstract if only statistics were examined.

Our model treated the four explanations as independent, under ceteris paribus conditions. The results, although strong and significant, raised some possible avenues for theorizing. For example, connections among the four explanations that could inform how developers make trade-offs between density overlap and dominance merits additional research. In cases that may indicate high expected growth and acceptance of a platform, developers may choose a new platform that is tightly coupled (that is, supported by few developers). We urge researchers to focus on interesting connections that may spawn further research. This phenomenon of network evolution is not unique to video games but exists in industries such as software, telecommunications, mobile Internet, and financial services. These industries could provide rich settings in which to refine theories and models of linkage formation and network evolution under fast-changing technologies.

Extensions of this study should address some of its limitations: (1) Our data set could be expanded to include game attributes—genre, functionality, links to movies, and so forth—to delve further into the intricacies of how video games and consoles coevolve. (2) Video game development could be studied to shed light on the ease of porting games to different platforms and on the restrictions on porting that console manufacturers may impose. (3) Future research could focus on additional links in this network, such as social connections (friendship ties among game developers), board interlocks, technical forums, and cross-equity or common-venture investment links. (4) The video game development process could also be studied to better specify the impact of time lags in decisions to launch games for platforms.

Our results offer some pointers for management practice. Interorganizational coordination of product architectures is not limited to video games but appears to operate in many other settings. Coordination of complementary offerings calls for a new competency focused on "network orchestration." This competency balances the inevitable tension between cooperation and competition across different types of resources over time and involves the formation and dissolution of multiple types of relationships. Since the locus of critical resources in such settings lies outside any one firm in a network (Gulati et al., 2000), network orchestration as a competency may be a critical driver of superior performance. Network orchestration is at a higher level than functional boundary spanning in purchasing, marketing, R&D, and operations. In periods of turbulent change, a network can be a source of strength (such as informational cues to threats of obsolescence) or a source of constraint (such as rigidity and inertia). Network orchestration as a competency entails managers' simultaneously focusing on the macro logic of network structure (how a portfolio of relationships is structured for resource access as a whole) and the micro logic of network processes (selection, cultivation, and dissolution of individual relationships) that contribute to maximal performance.

In conclusion, networks have become a popular topic for management research (Borgatti & Foster, 2003; Galaskiewicz & Wasserman, 1989; Gulati et al., 2000). Although alliances have been studied for some time, they have been mostly treated as dyadic connections—either static or dynamic. A network perspective is a powerful way to holistically understand the complex resource flows and dependencies that create performance differences. We made an initial exploration into the dynamics of network formation within system-based competition in one setting and found strong, encouraging results. We hope that this study will stimulate other researchers to better examine how networks of relationships confer competitive advantage over time. Such efforts will strengthen and broaden a network theory of strategic management.

REFERENCES

Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. ***Administrative Science Quarterly,*** 45: 425–455.

Albert, R., & Barabasi, A. 2002. Statistical mechanics of complex networks. ***Reviews of Modern Physics,*** 74: 47–97.

Arthur, W. B. 1989. Competing technologies, increasing returns, and lock-in by historical small events. ***Economic Journal,*** 99: 116–131.

Arthur, W. B., Durlauf, S. N., & Lane, D. A. 1997. The economy as a complex evolving system, II. ***Santa Fe Institute Studies in the Science of Complexity, Proceedings,*** vol. 27. Reading, MA: Addison-Wesley.

Baker, W. E., Faulkner, R. R., & Fisher, G. A. 1998. Hazards of the market: The continuity and dissolution of interorganizational market relationships. ***American Sociological Review,*** 63: 147–177.

Baldwin, C., & Clark, K. 2000. ***Design rules: The power of modularity.*** Cambridge, MA: MIT Press.

Barabasi, A. 2002. ***Linked: The new science of networks.*** Cambridge, MA: Perseus Publishing.

Batagelj, V., & Mrvar, A. 2003. Pajek—Analysis and visualization of large networks. In M. Jünger & P. Mutzel (Eds.), ***Graph drawing software (mathematics and visualization):*** 77–103. Berlin: Springer.

Baum, J. A. C., & Singh, J. V. 1994. Organizational niches and the dynamics of organizational mortality. ***American Journal of Sociology,*** 100: 346–380.

Baum, J. A. C., Shipilov, A. V., & Rowley, T. J. 2003. Where do small worlds come from? ***Industrial and Corporate Change,*** 12: 697–725.

Borgatti, S. P., & Foster, P. C. 2003. The network paradigm in organizational research: A review and typology. ***Journal of Management,*** 29: 991–1013.

Brandenberger, A. M., & Nalebuff, B. J. 1996. ***Co-opetition.*** New York: Currency Doubleday.

Bresnahan, T., & Greenstein, S. 1999. Technological competition and the structure of the computer industry. ***Journal of Industrial Economics,*** 47: 1–40.

Brown, S. L., & Eisenhardt, K. M. 1997. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. ***Administrative Science Quarterly,*** 42: 1–34.

Burt, R. S. 2000. Decay functions. ***Social Networks,*** 22: 1–28.

Burt, R. S., & Carlton, D. S. 1989. Another look at the network boundaries of American markets. ***American Journal of Sociology,*** 94: 723–753.

Davis, G. F., Yoo, M., & W. E. Baker 2003. The small world of the American corporate elite, 1982–2001. ***Strategic Organization,*** 1: 301–326.

Dyer, J. H., & Singh, H. 1998. The relational view: Cooperative strategy and sources of interorganizational competitive advantage. ***Academy of Management Review,*** 23: 660–679.

Evans, D. S., Nichols, A. L., & Schmalansee, R. 2001. An analysis of the government's economic case in US v. Microsoft. ***Antitrust Bulletin,*** 46(2): 163–252.

Frels, J. K., Shervani, R., & Srivastava, R. K. 2003. The integrated networks model: Explaining resource allocations in network markets. ***Journal of Marketing,*** 67(1): 29–45.

Galaskiewicz, J., & Wasserman, S. 1989. Mimetic processes within an interorganizational field: An empirical test. ***Administrative Science Quarterly,*** 34: 454–479.

Gallagher, S., & Park, S. H. 2002. Innovation and competition in standard-based industries: A historical analysis of the U.S. home video game market. ***IEEE Transactions on Engineering Management,*** 49: 67–82.

Garud, R., Kumaraswamy, A., & Langlois, R. (Eds.). 2003. ***Managing in the modular age.*** Malden, MA: Blackwell Publishing.

Gawer, A., & Cusumano, M. A. 2002. ***Platform leadership.*** Boston: Harvard Business School Press.

Granovetter, M. S. 1985. Economic action and social structure: The problem of embeddedness. ***American Journal of Sociology,*** 91: 481–510.

Greene, W. H. 2003. ***Econometric analysis.*** Upper Saddle River, NJ: Prentice-Hall.

Greene, W. H., & Hensher D. 1997. Multinomial logit and discrete choice models. In W. H. Greene, ***LIMDEP version 7.0 user's manual revised.*** New York: Planview.

Greve, H. R. 2000. Market niche entry decisions: Competition, learning, and strategy in Tokyo banking, 1894–1936. ***Academy of Management Journal,*** 43: 816–836.

Gulati, R. 1995. Social structure and alliance formation patterns: A longitudinal analysis. ***Administrative Science Quarterly,*** 40: 619–652.

Gulati, R., & Gargiulo, M. 1999. Where do interorganizational networks come from? ***American Journal of Sociology,*** 104: 1439–1493.

Gulati, R., Nohria, N., & Zaheer, A. 2000. Strategic networks. ***Strategic Management Journal,*** 21: 203–215.

Heide, J. B. 1994. Interorganizational governance in marketing channels. ***Journal of Marketing,*** 58(1): 71–85.

Katz, M. L., & Shapiro, C. 1994. Systems competition and network effects. ***Journal of Economic Perspectives,*** 8: 93–105.

Kogut, B., & Walker, G. 2001. The small world of Germany and the durability of national networks. ***American Sociological Review,*** 66: 317–335.

Liebowitz, S., & Margolis, S. E. 2001. ***Winners, losers & Microsoft: Competition and antitrust in high technology***. Oakland, CA: Independent Institute.

Marsden, P. V. 1981. Introducing influence and processes into a system of collective decisions. ***American Journal of Sociology,*** 86: 1203–1235.

McFadden, D. 1974. Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), ***Frontiers in econometrics:*** 105–142. New York: Academic.

McKenzie, R. 2000. ***Trust on trial: How the Microsoft case is reframing the rules of competition.*** Cambridge, MA: Belknap.

Nohria, N. 1992. Introduction: Is a network perspective a useful way of studying organizations? In N. Nohria & R. G. Eccles (Eds.), ***Networks and organizations: Structure, form, and action:*** 1–22. Boston: Harvard Business School Press.

Oliver, C. 1990. Determinants of interorganizational relationships: Integration and future directions. ***Academy of Management Review,*** 15: 241–265.

Podolny, J. M. 2001. Networks as the pipes and prisms of the market. ***American Journal of Sociology,*** 107: 33–60.

Podolny, J. M., & Stuart, T. E. 1995. A role-based ecology of technological change. ***American Journal of Sociology,*** 100: 1224–1260.

Powell, W. W., White, D. R., Koput, K. W., & Owen-Smith, J. 2004. Network dynamics and field evolution: The growth of interorganizational collaboration in life sciences. ***American Journal of Sociology:*** In press.

Salancik, G. 1995. WANTED: A good theory of organization. ***Administrative Science Quarterly,*** 45: 1–24.

Schilling, M. A. 2003. Technological leapfrogging: Lessons from the U.S. video game console industry. ***California Management Review,*** 45(3): 6–32.

Shankar, V., & Bayus, B. L. 2003. Network effects and competition: An empirical analysis of the home video game industry. ***Strategic Management Journal,*** 24: 375–384.

Shaprio, C., & Varian, H. 1998. ***Information rules: A strategic guide to the network economy.*** Cambridge, MA: Harvard University Press.

Sohn, M. 2001. Distance and cosine measures of niche overlap. ***Social Networks,*** 23: 141–165.

Stuart, T. 1999. A structural perspective on organizational innovation. ***Industrial and Corporate Change***, 8: 745–775.

Subramani, M. R. & Venkatraman, N. 2003. Safeguarding investments in asymmetric interorganizational relationships: Theory and evidence. ***Academy of Management Journal,*** 46: 46–62.

Uzzi, B. 1996. The sources and consequences of embeddedness for the economic performance of organizations: The network effect. ***American Sociological Review,*** 61: 674–698.

Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. ***Administrative Science Quarterly,*** 42: 35–67.

Uzzi, B., & Spiro, J. 2003. ***Do small worlds make big differences? Artist networks and the success of Broadway musicals, 1945–1989.*** Working paper, Northwestern University, Evanston, IL.

Van de Bunt, G. G., Van Duijn, M., & Snijders, T. A. B. 1999. Friendship networks through time: An actor-oriented dynamic statistical network model. ***Computational & Mathematical Organization Theory,*** 5(2): 167–199.

Van de Ven, A. H. 1976. On the nature, formation and maintenance of relations among organizations. ***Academy of Management Review,*** 1: 24–36.

Watts, D. 2003. ***Six degrees: The science of a connected age.*** New York: W. W. Norton.

Williamson, O. E. 1991. Strategizing, economizing, and economic organization. ***Strategic Management Journal,*** 12: 75–94.

APPENDIX

Assembling and Validating the Data Set

We extracted new U.S. console title, platform, developer, and release date information from three Web sites (www.gamefaqs.com, www.gamespot.com, and www.ign.com). The data were then cross-validated with the lists provided on console manufacturers' Web sites. Nonprogrammable single games (e.g., "Mattel Football") and PC platform titles were discarded. Also discarded were games for portable handheld equipment, such as Nintendo's Gameboy Color and Sega's GameGear, and titles released only for Japanese and other non-U.S. markets.

Semantic variations were resolved and redundancies removed for all developers, platforms, titles, and release dates using (1) the Web sites of platform manufacturers (e.g., www.playstation.com) and game developers (e.g., www.ea.com), (2) independent video game databases (e.g., gamespot.com, ign.com, and vgmuseum.com), and (3) on-line retailer databases (e.g., bestbuy.com and amazon.com). We collected additional information from popular video game publications. A complete list of sources used to validate the data is available on request.

An observation was used if we got consistent data on it from at least three sources. We also discarded titles that were rereleased for the same platform in a subsequent year as well as titles that consisted of bundles of previous releases. However, existing titles released for new platforms were included because games rarely run on more than one platform. Also discarded were titles that could operate on multiple consoles through the use of game emulators. Although some recent platforms by a single manufacturer (e.g., Sony PlayStation 2) will accept games written for a previous platform (e.g., Sony PlayStation), these games are rarely rereleased. If they were rereleased, we deleted data on them from the database.

Finally, we obtained yearly data for the top 100 titles for each year (1995–2002) from the NPD Group. Each of the 800 observations represents a unique platform, developer, title, units shipped, unit share, and year combination. The top 100 titles on average accounted for 68 percent of sales, thus allowing for reasonable validation. We used these data to recalculate some variables, counting units shipped instead of number of new titles.

N. Venkatraman *(venkat@bu.edu)* is the David J. McGrath Jr. Professor of Management at the Boston University School of Management. He received his doctorate from the University of Pittsburgh. His research interests are at the intersection of strategic management and information technology, with an emphasis on competing through networks.

Chi-Hyon Lee *(chihyon@bu.edu)* is an assistant research professor at the Boston University School of Management. He received his doctorate from Boston University. His research interests are software ecosystem dynamics, alliances and partnerships, and network analytics and methods.