MANAGEMENT SCIENCE
Vol. 55, No. 4, April 2009, pp. 635–649
ISSN 0025-1909 | EISSN 1526-5501 | 09 | 5504 | 0635

informs

DOI 10.1287/mnsc.1080.0968

Alliance Structure and the Scope of Knowledge Transfer: Evidence from U.S.-Japan Agreements

Joanne Oxley
Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada,
oxley@rotman.utoronto.ca

Tetsuo Wada
Faculty of Economics, Gakushuin University, Tokyo 171-8588, Japan,
tetsuo.wada@gakushuin.ac.jp

Prior research suggests that equity joint ventures (JVs) are particularly effective vehicles for accessing complex technology. Different schools of thought have emphasized different reasons why joint ventures might support greater knowledge transfer than "bare" license agreements: incentive alignment, organizational embeddedness, and enhanced administrative controls. We probe and refine these theoretical perspectives, drawing out implications of the different theories for the extent and speed of alliance-related knowledge transfer, as well as for knowledge "leakage" in areas not directly related to alliance activities. Using a proprietary data set derived from regulatory filings with the Japanese government we test these implications in our empirical analysis of U.S.-Japan agreements. The picture that emerges from the analysis is one of particularly intense but contained knowledge transfer in equity joint ventures, relative to bare license agreements: knowledge transfers directly related to the alliance activity are enhanced in the JV, and the speed of integration into the Japanese firm's subsequent innovations also increases. In marked contrast, leakage of unrelated technology is significantly reduced. These findings suggest that administrative structures that reduce technology leakage are a key feature of the equity joint venture, a result that is inconsistent with a "pure" knowledge-based perspective on alliances.

Key words: patent licensing; joint ventures; U.S.-Japan alliances; knowledge transfer
History: Received January 19, 2006; accepted November 6, 2008, by Scott Shane, technological innovation, product development, and entrepreneurship. Published online in *Articles in Advance* January 28, 2009.

1. Introduction

Effective knowledge management is a key issue underlying competitive advantage, particularly in high-technology industries. To build sustainable competitive advantage, firms must integrate and apply knowledge created internally with knowledge acquired from external sources. As external knowledge sources are increasingly global and diverse, interfirm alliances have become popular mechanisms by which firms access external knowledge. Paralleling this rise in popularity, academic interest in alliances as knowledge acquisition mechanisms has also grown, generating a significant body of empirical research over the past 15 years. It is now generally accepted that participation in an alliance produces measurable changes in the technological resources of partner firms and that equity joint ventures (JVs) are particularly effective knowledge transfer vehicles as compared with more arms-length contractual agreements (Mowery et al. 1996, Gomes-Casseres et al. 2006).

Various theoretical perspectives have been invoked to explain why the equity joint venture structure might facilitate greater knowledge transfer: transaction cost economists focus on the incentive alignment associated with shared equity (e.g., Pisano 1989, Oxley 1997); researchers approaching alliances from a knowledge-based perspective argue instead that joint ventures create organizationally embedded communities that facilitate tacit knowledge transfer (e.g., Kogut 1988, Kogut and Zander 1992); meanwhile, scholars in international business point to the enhanced control over technology and strategy that comes with equity ownership (e.g., Geringer and Hebert 1989). All of these perspectives are consistent with the idea that the joint venture structure encourages knowledge sharing to facilitate alliance activities, even as they highlight quite different organizational mechanisms.

In this paper we probe and refine these three schools of thought, applying them to an empirical analysis of technology acquisition by Japanese firms through alliances with U.S. companies. We draw out implications of the different theories for the extent and speed of alliance-related knowledge transfer, as well as for knowledge "leakage" in areas not directly related to alliance activities. Extant empirical research on the scope of knowledge transfer is quite sparse, reflecting difficulties in accessing adequate data and devising measures of knowledge flows in different areas,

something that bedevils all empirical work in knowledge management (Almeida et al. 2002). In this study, however, we are fortunate to have access to a proprietary data set uniquely suited to addressing questions related to the scope of knowledge transfer. The data set is derived from regulatory filings with the Japanese government and was obtained through the National Institute for Science and Technology Policy (NISTEP); it comprises information on all patent license agreements initiated between U.S. and Japanese firms during the period 1988–1991. Because we have information about the specific patents covered by each agreement, we are able to distinguish knowledge transfer in the technological domain of the licensed patents from transfers in other, unrelated domains. Linking these data with information on joint venture agreements in effect at the beginning of the sample period allows us to compare the scope of related and unrelated knowledge transfers observed in "bare" unilateral or cross-license agreements with those that take place within the context of an ongoing joint venture.

Our empirical strategy relies on changes in the extent to which Japanese licensees' patent applications cite the patents of their U.S. alliance partners in the periods before and after establishment of the license agreements. Following prior literature, we interpret patent citation as an indicator of knowledge building (Almeida 1996); we also examine changes in the "citation lag" (i.e., the average age of patents cited in new patent applications) as a measure of the speed of knowledge integration (Sampat et al. 2003, Fabrizio 2007).

Our empirical analyses yield robust results that further our understanding of alliances as technology acquisition vehicles. Extending the findings of prior research we show that Japanese JV partners more frequently cite their U.S. licensor's patents in the technology domain of the license (controlling for prelicense citations) and that the speed of knowledge integration also increases. In a novel departure from previous research, we also show that joint ventures *reduce* knowledge flows in technology areas unrelated to alliance activities when compared with bare license agreements. This finding highlights the importance of administrative control as a key distinguishing feature of the equity joint venture structure and suggests that enhanced cooperation in alliance-related areas is predicated on the ability of alliance partners to control unintended leakage in areas beyond the scope of the alliance.

The remainder of this paper proceeds as follows. In §2, we situate our research in the prior literature, develop the theoretical arguments, and present our hypotheses. In §3, we describe the data and methods used in our empirical analysis. Presentation and discussion of the results can be found in §4, followed by robustness checks and exploration of alternative explanations in §5. In §6, we discuss implications of our results for future research.

2. Theory and Hypotheses

The logic underlying the assertion that equity joint ventures support greater knowledge transfer than contractual alliances is multifaceted. Transaction cost economists suggest that one important distinguishing feature of different knowledge transfer alliances is the degree of *incentive alignment*. They suggest that alliance types may be conceptualized as lying along a market-hierarchy continuum (Oxley 1997), from the most arms-length unilateral technology license agreements through cross licensing or other contract-based technology-sharing agreements to equity joint ventures, which lie closer to the governance properties of internal organization or hierarchy. Moving from a unilateral to a cross-licensing agreement increases incentive alignment because of the reciprocal nature of the agreement: each partner risks forgoing access to the other's technology should they behave opportunistically. Incentive alignment increases further in an equity joint venture, where ownership is shared by the partner firms, because any reduction in returns due to opportunism is also shared. Equity thus promotes cooperation and knowledge sharing within the domain of the alliance.

Looking at equity joint ventures from a knowledge-based perspective, other researchers have argued that the key difference between equity joint ventures and contract-based alliances (such as license agreements) lies in the fact that joint ventures are "organizationally embedded" (Kogut 1988). In this view, organizational embeddedness facilitates the development of a shared identity and common language. This in turn encourages informal learning behaviors in the form of spontaneous interaction and knowledge sharing during the course of alliance activities. Such informal learning behavior is thought to be particularly important for promoting a shared understanding and facilitating tacit knowledge flows (Kogut 1988, Elkjaer 2003), so enhancing alliance outcomes. The knowledge-based perspective thus emphasizes the *coordination* benefits of the joint venture structure and suggests that the joint venture's organizational embeddedness reduces constraints on broad knowledge sharing: potential opportunism by alliance partners is seen as a minor consideration in alliance design, relative to the organizational complexity of the knowledge transfer process.

Finally, contributors to the international business literature on joint ventures and alliances have emphasized the point that equity joint ventures also embody

complex administrative structures that enhance *control* over technology flows and other transfers from one partner to the other in an alliance (e.g., Geringer and Hebert 1989, Balakrishnan and Koza 1993).[1] From a control perspective, perhaps the most salient feature of an equity joint venture is the legal entity created, which is jointly owned but separate from either parent. The resulting legal and administrative boundary between alliance activities and other operations of the partner firms enhances the range of control mechanisms available to the firm supplying technology to the venture relative to those in a license agreement directly connecting the two firms.

The additional controls available to joint venture partners include an increased ability to enforce employee conduct rules and job responsibilities, as well as an ability to manage the mobility of engineers working on alliance activities, including those who were previously employed by the partner firm; such policies are very difficult if not impossible to enforce in a license or other contractual agreement (Liebeskind 1996). Control of employee mobility is particularly important because it decouples informal learning through unstructured interaction *within* the joint venture from more formal knowledge transfer mechanisms, such as programmed events, meetings, visits, etc., that engage the parent firms' organizations. As noted above, informal learning behavior is important for transfer of tacit knowledge, but the unstructured nature of these spontaneous interactions may also increase the risk of unintended knowledge leakage if, for example, engineers working on alliance activities interact freely with peers in their respective home organizations (Janowicz-Panjaitan and Noorderhaven 2008). The formalized boundary-spanning channels between a joint venture and the parent companies, in contrast, allow for more targeted and contained knowledge transfer focusing on technology and other knowledge inputs necessary to achieve the specific goals of the alliance. In this way, knowledge flows into and out of the joint venture can be limited to alliance-relevant knowledge without undermining the effectiveness of knowledge transfer and learning within the joint venture itself.

Overlaid on the joint venture's legal and organizational structure are a range of additional administrative mechanisms that also facilitate control of knowledge flows into and out of the venture. First among these is the joint management board, which provides a direct communication link with senior managers of the parent companies and also provides for greater oversight and monitoring of alliance employees and partner firm activities (Geringer and Hebert 1989, Yan and Gray 1994).[2] In addition, JV partners have enhanced auditing rights over both financial and operational activities related to the venture that are rarely available to licensors (Osborn and Baughn 1990). In a survey examining contractual provisions used in alliances in the German telecommunication industry, for example, Reuer et al. (2006) show that equity joint ventures commonly feature provisions related to reporting and auditing rights that are usually absent from nonequity agreements.[3] In the context of a technology transfer alliance, the enhanced control and monitoring rights that accrue to the technology provider can be used to bound knowledge flows in an attempt to reduce potentially harmful leakage of technology and intellectual property in areas that are not directly relevant to the successful execution of alliance activities (Oxley and Sampson 2004).

Before we turn to hypothesis development, it is useful to situate the theoretical arguments in the particular context of our study: in-licensing of U.S. technology by Japanese firms in the late 1980s to early 1990s. During this period, Japanese firms licensed a broad range of technologies from foreign (predominantly U.S.) companies as they approached and subsequently pushed the technological frontier in many industries. This in-licensing activity took a variety of forms. In some cases, Japanese firms simply acquired the right to use a U.S. firm's technology in its own products or as part of an OEM manufacturing agreement with the U.S. patent holder to facilitate Japan-based production. These bare unilateral licenses granted usage rights to one or more patents and also provided access to technological know-how (e.g., blueprints, training, start-up consulting), most often facilitated by interactions between the engineers of the two companies to ensure the required transfer of tacit knowledge associated with the technology. In other cases, a particular in-license may be part of a broader cross-license or technology sharing agreement between the U.S. and Japanese firms. Again, there would usually be active knowledge transfer and technological cooperation beyond provision of blueprints, etc., and such agreements would typically include some provision for access to follow-on inventions by one or both parties.[4] Finally, license

[1] Transaction cost economics (TCE) also is concerned with the administrative controls available in different alliance forms (see Oxley 1997), but this area of TCE theory is relatively underdeveloped.

[2] In their case study of U.S.-China joint ventures, for example, Yan and Gray (1994, p. 1493) find that "nominations of members of a venture's board of directors and general manager are important control mechanisms" and that the "board of directors is empowered to discuss and take action on all fundamental issues concerning the venture."

[3] Forty-four percent of the contracts governing equity joint ventures in the Reuer et al. (2006) sample stipulate specific auditing rights, and 78% stipulate rights to reports of relevant transactions. This compares with 15% and 41%, respectively, for similar auditing and reporting provisions in nonequity alliances.

agreements (whether unilateral or cross-licenses) may occur within the context of an equity-based joint venture agreement linking the Japanese and U.S. firms.[5] As with the case of bare technology licenses, such ventures were typically set up to manufacture products in Japan based on U.S. technology (either alone or in combination with Japanese technology). In many cases these joint ventures were operationally distinct from existing operations of the Japanese firm; at the same time, they provided access to technology that could have application in other areas of the firm's business.

The theoretical arguments above, regarding enhanced incentive alignment, coordination, and control in equity joint ventures relative to contractual alliances, imply that alliance partners will cooperate more intensely in alliance-related areas in joint ventures than in bare licenses and that unrelated knowledge flows to the partner will also be more heavily controlled in this case. Prior case studies of Japanese-U.S. alliances provide some evidence of this behavior. In Fuji Xerox, for example, a major focus of the joint venture agreement and organizational structure adopted was the prevention of uncontrolled knowledge flows to the Japanese partner's operations beyond those directly related to alliance operations (Gomes-Casseres 1991).[6] And in a JV linking Caterpillar Inc. and Mitsubishi Heavy Industries, Caterpillar rejected an initial proposal by the Ministry of International Trade and Industry that it license its technology to a Japanese company and held out instead for an equity joint venture agreement, reportedly because of concerns that a license agreement could potentially lead to the development of a broad-based competitor (Allen 1991).

This anecdotal evidence is consistent with the notion that the joint venture structure gives a technology provider greater control over unintended leakage to the alliance partner. One might still wonder, however, why the partners could not replicate these control features in a less costly license agreement.[7] In Japan, the historical development of the legal system, along with apparent cultural avoidance of legal action, means that reliance on detailed contingent contracts is rare (Gilson and Roe 1993). At the same time, Japanese law confers significant managerial control rights on minority shareholders, including veto rights over board decisions (Ahmadjian and Oxley 2006). This right ensures that in an equity joint venture linking a U.S. and Japanese firm, the U.S. firm retains substantial control over the activities of the venture, including interactions with the Japanese partner firm.[8] Given the limitations of contracts in this context, such control is most certainly not available in a bare license agreement.

We can now build on the arguments above to generate empirical implications for technology in-licensing in Japan and the extent to which a licensee integrates and leverages the technology of its U.S. partner in the years after the focal licensing activity. If knowledge transfer is primarily in areas directly related to the alliance activities per se, we would expect subsequent technology developed by the Japanese firm to build on technologies incorporated in the license or on closely related technologies. Each of the different theoretical perspectives discussed above supports the contention that the joint venture structure promotes enhanced knowledge transfer in areas directly related to alliance activities, so we would expect alliance-related technology building by the Japanese firm to be higher when the focal license takes place in a joint venture. When it comes to the relative knowledge transfer characteristics of unilateral and cross-license agreements, however, only transaction cost economics speaks directly to the issue: the incentive alignment arguments of TCE imply that cross-licenses will be intermediate in their knowledge transfer features—there is greater incentive alignment among partners in a cross-licensing agreement than in a unilateral license, but the incentive alignment is not as strong as in an equity joint venture. These arguments lead to our first set of hypotheses:

Hypothesis 1A (H1A). *After the initiation of a technology license, alliance-related technology building by the*

[4] It is also possible that some license agreements are initiated in response to actual or threatened litigation and do not involve knowledge transfer per se. Although we cannot distinguish in our empirical analysis between licenses designed for technology sharing and licenses intended only to establish "litigation-free zones," inspection of the data suggests that the latter type do not represent a significant fraction of our sample licenses because many of the contracts include forward-looking provisions such as know-how transfer clauses, etc.

[5] For equity joint ventures the patent license may be granted to the joint venture itself or directly to the Japanese parent company.

[6] "According to Fuji Photo Film's agreement with Xerox, the company (FPF)... could collect information from Fuji Xerox, but it could not use it in its own operations" (Gomes-Casseres 1991, p. 5). The intensity of cooperation *within* the alliance domain is evidenced by the great success of Fuji Xerox in becoming a leader in low- to mid-range copiers.

[7] Interestingly, prior literature on alliances provides little direct insight into this question: despite its centrality to theories of alliance governance, the assumption that the control features of an equity-based venture are unavailable to partners in a contractual alliance is largely implicit in prior research (e.g., Oxley 1997, Almeida et al. 2002, Santoro and McGill 2005). See the concluding section for more on this issue.

[8] The ownership threshold for control rights under Japanese law is 33.4%. It would be highly unusual for a foreign partner to hold less than this share in a bilateral U.S.-Japan joint venture. In our sample, for example, the minimum equity holding of the U.S. partner is 45% (see Footnote 11 for more details on ownership shares in our sample of alliances).

Japanese firm is higher if the licensee and licensor are equity joint venture partners.

Hypothesis 1B (H1B). *After the initiation of a technology license, alliance-related technology building by the Japanese firm is higher if the license is part of a cross-license agreement as compared with a unilateral license; alliance-related technology building will not be as high as for joint venture partners*.

Incentive alignment speaks primarily to the extent or domain of knowledge transfers. Enhanced coordination through organizational embeddedness, on the other hand, speaks to the potential for increased *speed* of technology transfer in equity joint ventures. Organizational embeddedness is associated with "high bandwidth" communication channels, wherein teams are colocated, allowing them to work closely together, creating common understandings of how to transfer and transform knowledge (Tallman 2003), enhancing problem solving and facilitating the transfer of tacit knowledge (Heiman and Nickerson 2004). If indeed joint ventures are better suited to the development of such high-bandwidth communications channels, then the speed of knowledge integration and technology building by the Japanese licensee should be higher in joint ventures relative to bare license agreements (Zander and Kogut 1995). We thus offer the following hypothesis:

Hypothesis 2A (H2A). *After the initiation of a technology license, the speed of knowledge integration and technology building by the Japanese licensee will be higher if the licensee and licensor are equity joint venture partners*.

If we restricted our attention to the knowledge-based view of the firm we would offer no corollary to this hypothesis for cross-license agreements: there is little basis for arguing that cross-licenses embody greater organizational embeddedness than unilateral licenses because both are contractual agreements and not organizationally instantiated (Kogut 1988). However, adoption of high-bandwidth communication channels and other knowledge management practices is not conceptually bound to any particular organizational form (Foss 1996). Indeed, transaction cost economists argue that incentive alignment is an essential underpinning of partners' willingness to adopt high-bandwidth channels (Oxley 2004). If this is the case, then the elaboration of knowledge management practices (and therefore the speed of knowledge transfer) should parallel the market-hierarchy continuum of alliance forms, with cross-licenses lying intermediate between unilateral licenses and equity joint ventures. Thus, we have:

Hypothesis 2B (H2B). *After the initiation of a technology license, the speed of knowledge integration and technology building by the Japanese licensee will be higher if the license is part of a cross-license agreement as compared with a unilateral license; the speed of knowledge integration will nonetheless be slower than for joint venture partners*.

As emphasized by proponents of the control perspective, equity joint ventures also embody stronger rights (and mechanisms) for control of knowledge flows. In the context of a technology transfer alliance, we predict that the technology provider will use this enhanced control and monitoring ability to bound knowledge flows in an attempt to reduce potentially harmful leakage of valuable intellectual property in areas that are not directly relevant to the successful execution of alliance activities. This leads us to our final hypothesis:

Hypothesis 3 (H3). *After the initiation of a technology license, technology building by the Japanese firm in areas unrelated to alliance activities will be lower if the licensee and licensor are equity joint venture partners*.

This last hypothesis deserves additional comment because it reflects a "strong-form" version of the control argument relative to the coordinating structures of joint ventures and places significant reliance on the "ceteris paribus" condition that naturally applies to all of the hypotheses. Because JVs feature greater organizational embeddedness and high-bandwidth knowledge transfer mechanisms one might expect knowledge flows to be greater in *all* areas, whether related or unrelated to the alliance activities. Similarly, one could argue that the incentive alignment achieved within equity JVs encourages the partners to do "whatever it takes" to achieve the common goal of the venture, which could include more broad knowledge sharing that does not constitute leakage per se. If joint ventures are inherently broader in their knowledge scope then, even if control mechanisms reduce leakage and unrelated knowledge flows, the net effect may be ambiguous. H3 thus represents a strong test of the control hypothesis, raising the hurdle for significant findings in our empirical analysis, to which we now turn.

3. Data and Methods

Particularly detailed data on Japanese technology in-licensing from the United States exist because Japanese firms were required to report technology importing contracts (patent and know-how licensing) under Article 29 of the former Foreign Exchange Law of Japan, in effect until 1998. We were able to gain confidential access to data from the National Institute for Science and Technology Policy (NISTEP) on the content and terms of all patent license agreements between Japanese and U.S. firms during the period 1988–1991. These data encompass 536 licensor-licensee firm pairs, with 602 patent licensing contracts, involving

1,786 patents.[9] More than 80% of these contracts are in the electronics and machinery sectors. The contracts are a mixture of unilateral and cross-license agreements, although we have information on the content and terms of the in-license only.[10] Thus, all of the licensors in our empirical analysis are U.S. firms—257 firms in all—and all of the licensees are Japanese firms—167 in total.

We match the licensing data with information on joint venture agreements in effect at the beginning of the sample period, documented in the *Toyo Keizai Directory of Foreign Subsidiaries in Japan* and *Toyo Keizai Directory of Foreign Subsidiaries by Japanese Firms*.[11] Of the 536 firm dyads in our sample, 27 firm pairs have equity joint ventures in effect at the beginning of the sample period. The data are also matched with data on U.S. patenting activities of the U.S. licensors and Japanese licensees from 1975 to 2002 drawn from the National Bureau of Economic Research (NBER) patent data file (Hall et al. 2001) and the Micropatent database.[12]

Following prior research we focus on technology developments embodied in patents and operationalize the scope of knowledge transfer and technology building using U.S. patent citation data (e.g., Almeida 1996, Mowery et al. 1996, Almeida et al. 2002).[13] Although the use of patent data is subject to various criticisms, patent citations remain the most well accepted indicator of the technological lineage of new patents. Every U.S. patent application must include citations to all existing patents that constitute relevant "prior art"—i.e., the patented portion of the existing technology pool from which the new innovation draws. Consequently, as a Japanese firm acquires technological knowledge from its U.S. alliance partner we should see a higher rate of citation of the U.S. firm's patents in new patents applied for by the Japanese firm. Patent citations also give an indication of the technological domain of the knowledge transferred in an alliance. Patents are classified by the patent examiner into technology "classes" that represent different technological areas. These technology classes have been used in prior empirical research as an indication of a firm's areas of technological expertise (e.g., Jaffe 1986); we thus use information on the technology classes covered by the patents held by each alliance partner firm to construct measures of the extent to which a firm builds on its alliance partner's knowledge in different technological areas.[14] More specifically, we distinguish between citations to patents that belong to the same patent classes as the licensed patents—*related* citations—and citations to patents that belong to patent classes not covered by the licensed patents—*unrelated* citations.

The unit of analysis in our empirical study is the licensor-licensee dyad, and the scope of knowledge transfer is captured using a variety of patent citation measures, which form the dependent variables in our regression analysis:

- RELATED POSTLICENSE CITES$_{ij}$ = number of cites to any licensor$_j$ patent (1975–2002) in licensee$_i$'s patents applied for after 1992,[15] where the cited patent is in the same technology class(es) as the licensed patent(s).

This measure captures the extent to which the licensee builds on licensor technology in areas related to those of the focal licensing agreement, after initiation of the license;

[9] Internal transactions—i.e., transfers from a U.S.-based subsidiary to a Japanese-based parent company—are reported as such in the NISTEP data. We deleted these to avoid contaminating the analysis with self-citations.

[10] NISTEP requires firms to indicate whether a particular license is part of a cross-license agreement, but we do not have information on what patents are out-licensed by the Japanese firm to its U.S. counterpart.

[11] JVs listed in these directories all involve the establishment of a separate entity jointly owned by the parent companies. More than half of the JVs in our sample have exactly 50:50 ownership; ownership range for the U.S. companies in the other ventures ranges from 45% to 75%. We found no systematic differences in the characteristics or effects of these minority- or majority-owned ventures as compared with the equal partnerships.

[12] Following prior research we build patent portfolios at the corporate level, i.e., including patents assigned to a parent company plus all of its subsidiaries. The criterion used to define subsidiaries for Japanese firms is whether Toyo Keizai's 1990 *Nihon no Kigyo Grupu* designated a subsidiary firm as "consolidated." Only one JV in our sample itself held patents, and none of these patents generates any relevant citations that might contaminate our analysis.

[13] Some of the patent licenses in our sample cover Japanese patents held by the U.S. firm (439 patents in all). In these cases we use the Organisation for Economic Co-operation and Development's "triadic patent families" (Dernis and Khan 2004) to classify citations into related and unrelated categories. Patent citation counts in the empirical analysis are restricted to U.S. patents.

[14] Note that we use the 118 technology classes defined in the International Patent Classification System to construct our measures. This classification is more fine-grained than the U.S. Patent and Trademark Office classes used in Jaffe (1986) and related studies and avoids some of the problems associated with arbitrary aggregation.

[15] Because the latest license in our sample was initiated in December 1991, we can safely assume that patent applications submitted after 1992 reflect "postlicense" knowledge stocks of the licensee. An alternative approach would involve dating postlicense patent citations from the year after the initiation of the focal license agreement (e.g., using patent applications dating from 1989 or later for licenses initiated in 1988). The advantage of our approach is that we avoid the problem of comparing citation rates from patent portfolios of different vintages, a particularly important issue given documented changes in citation rates over time (Jaffe and Trajtenberg 2002) and the progression of Japanese firms toward the technological frontier during this period. We control for the possibility that earlier licenses yield greater citation rates in the post-1992 period (due to citation lags) by including year dummies in our regressions.

• UNRELATED POSTLICENSE CITES$_{ij}$ = number of cites to any licensor$_j$ patent (1975–2002) in licensee$_i$'s patents applied for after 1992, where the cited patent is not in any of the technology classes covered by the licensed patent(s).

This measure indicates the extent to which the licensee builds on licensor technology in areas unrelated to those of the focal license, after initiation of the license;

• POSTLICENSE BREADTH$_{ij}$=UNRELATED POSTLICENSE CITES$_{ij}$/TOTAL POSTLICENSE CITES$_{ij}$, where

• TOTAL POSTLICENSE CITES$_{ij}$ = number of cites to any licensor$_j$ patent (1975–2002) in licensee$_i$'s patents applied for after 1992.

POSTLICENSE BREADTH thus captures the extent to which changes in technology building under different alliance structures reflect symmetric changes in both related and unrelated technology areas or whether there are significant changes in the breadth of knowledge transfers.

We also use patent data for our analysis of the impact of enhanced coordination on the speed with which the licensor's technology is incorporated into the licensee's subsequent inventions. Again following prior research (Sampat et al. 2003, Fabrizio 2007), we use average citation lag, i.e., the lag between the time that a licensor's patents are granted and the time that they are cited in new patent applications by the licensee, as an indicator of the speed of technology transfer. As with the citation count measures, we distinguish between citations to related and unrelated patents:

• POSTLICENSE CITATION LAG (RELATED)$_{ij}$ = the average age of the licensor$_j$ patents (1975–2002) cited in licensee$_i$'s patents applied for after 1992,[16] where the cited patent is in one of the technology classes covered by the licensed patent(s);

• POSTLICENSE CITATION LAG (UNRELATED)$_{ij}$ = the average age of the licensor$_j$ patents (1975–2002) cited in licensee$_i$'s patents applied for after 1992, where the cited patent is not in any of the technology classes covered by the licensed patent(s).

The independent variables of primary interest in the empirical analysis are related to the form of the alliance. As discussed earlier, in our data we are able to distinguish between unilateral and cross-licensing agreements, and we can also observe whether the license took place within the context of an equity joint venture. We thus adopt the following measures of alliance form:

• CROSS LICENSE$_{ij}$ = 1 if the focal alliance involves a cross-licensing agreement, 0 if unilateral license only;

• JV$_{ij}$ = 1 if the licensor and licensee are partners in an equity joint venture at the time the license is initiated; 0 otherwise.

We also include a variety of control variables that capture firm or patent characteristics that may affect the propensity of a licensee to cite a licensor's patents in the postlicense period but that are unrelated to the alliance activity per se. First, it has been established in prior research that citation patterns exhibit strong serial correlation (Jaffe and Trajtenberg 2002) so that a firm whose previous patents tended to cite a particular set of "prior art" is likely to cite the same prior art in future patent applications. We therefore include prelicense citations as a control variable in the regressions using either related or unrelated patent citations (or lags) or prelicense breadth, as appropriate:

• RELATED/UNRELATED/TOTAL PRELICENSE CITES$_{ij}$ = number of citations to licensor$_j$ patents in licensee$_i$'s patents applied for before 1988;[17]

• PRELICENSE BREADTH$_{ij}$ = UNRELATED PRELICENSE CITES$_{ij}$/TOTAL PRELICENSE CITES$_{ij}$;

• PRELICENSE CITATION LAG (RELATED/UNRELATED)$_{ij}$ = the average age of the licensor$_j$ patents (1975–2002) cited in licensee$_i$'s patents applied for before 1988.

It is also well documented that firms with highly similar or overlapping technological portfolios are likely to cite each other's patents (Jaffe and Trajtenberg 2002). Thus, to ensure that high citation rates observed among alliance partners are not driven simply by preexisting technological relatedness rather than by alliance activities per se we include a measure of prelicense technological overlap between the partner firms. Following prior research, we use the measure of technological overlap or proximity developed by Jaffe (1986) based on the angular separation of the patent class distribution vectors of the partner firms.

• OVERLAP$_{ij}$ = Jaffe measure of technology overlap between licensee$_i$ and licensor$_j$ based on patents applied for in the 10 years prior to the sample period.[18]

[16] Again we use a fixed date to mark the beginning of the postlicense period, rather than keying the date off the individual license agreement. We include year dummies in our regressions to account for the different initiation dates of the sample alliances, and we also reran the analysis on individual year samples to evaluate the robustness of our results. The direction of the effects is identical across all individual years, and they are of the same order of magnitude.

[17] Because the first license in our sample was granted in 1988, patent applications submitted prior to 1988 clearly reflect "prelicense" knowledge stocks of the licensee.

[18] The precise construction of this variable is as follows: let F_i^s represent the number of patents assigned to firm i in class s. $F_i = (F_i^1, \ldots F_i^s)$ is then the patent class distribution vector for firm I and OVERLAP$_{ij}$ is calculated as $F_i\ F_j'/\sqrt{(F_i\ F_i')\ (F_j\ F_j')}$. See Jaffe (1986) for

Another relevant factor is the "importance" of the licensor's patents (Jaffe and Trajtenberg 2002) because important patents that receive a high number of citations overall are more likely to be cited by any given firm, ceteris paribus:

• PATENT IMPORTANCE$_j$ = total citations (divided by 1,000) in all post-1992 U.S. patent applications (by any firm) to licensor$_j$'s patents.

In a similar manner, by simple arithmetic, the total number of citations to prior art by a particular licensee will be a function of the total number of patents that the firm applies for in a given period:

• LICENSEE$_i$ PATENTS = total number of patents (divided by 1,000) granted to licensee$_i$ for which the application date is after 1992.

Controlling for licensee patenting activity also allows us to control (roughly) for potential heterogeneity in the capabilities of the Japanese licensees: a licensee with extensive patent holdings is likely to be a relatively "capable" firm with higher overall absorptive capacity. To further account for potential heterogeneity in licensee quality, we include measures of licensee size, profitability, and previous joint venture experience:

• LICENSEE$_i$ SIZE = Total equity capital (millions of yen) as reported in Toyo Keizai's *Nihon no Kigyo Grupu* (directory of Japanese firms and domestic subsidiaries), 1990;

• LICENSEE$_i$ PROFITS = After-tax profits (millions of yen) as reported in Toyo Keizai's *Nihon no Kigyo Grupu*, 1990;

• LICENSEE$_i$ JV EXPERIENCE = number of joint venture agreements established with other firms in the 10 years prior to the sample period.

Another potential concern is the possibility of spuriously high "unrelated" citations in cases where the licensor or licensee patent portfolios are particularly diverse. We therefore include measures of the concentration of the licensor and licensee patent holdings:

• LICENSOR$_j$ TECHNOLOGY CONCENTRATION = Herfindahl index based on patent classes in all licensor$_j$'s patents applied for before 1988;

• LICENSEE$_i$ TECHNOLOGY CONCENTRATION = Herfindahl index based on patent classes in all licensee$_i$'s patents applied for before 1988.

Finally, we include two sets of dummy variables: one set indicating the year in which the first license agreement between a particular licensee and licensor appears in our data (1988–1991) and one set of 12 dummy variables identifying the general technological sector of the alliance—textiles, chemicals, pharmaceuticals, new materials, metals, machinery, electronics, telecommunications, transportation, instruments, other manufacturing, and services. These variables, derived from the NISTEP filings, control for possible differences in patenting behavior across technological domains.

We test our hypotheses by examining the relationship between each of our measures of knowledge transfer and this set of independent variables. Because of the nature of patent citation data (count data with high dispersion), negative binomial regression analysis is most appropriate for our citation analysis; the analyses of postlicense breadth and citation lags use ordinary least squares regression. In all cases, we account for the fact that some licensors have alliances with more than one licensee by estimating robust standard errors, clustering on licensors. Because H1 and H3 are operationalized using citation counts whereas H2 is operationalized via citation lags, we first present results of regressions relevant to H1 and H3, followed by H2.

4. Empirical Results

Table 1 shows summary statistics for the complete sample of dyads, as well as for the 27 licensor-licensee pairs with an equity joint venture in effect at the beginning of the sample period. Note that the subsample of firm dyads with joint ventures are quite similar to the overall population of licensor-licensee pairs, with one notable exception: despite similar levels of overall technological overlap and slightly fewer licensee patents, Japanese joint venture partners cite their U.S. licensor at a much higher rate than non-JV partners, whether in related or unrelated areas, both before and after initiation of the license. A likely explanation is that the Japanese firms in the JV dyads have already acquired significant technological knowledge from their joint venture partner at the beginning of the sample period, and this is reflected in the higher citation rate.

The results of our first set of analyses are reported in Table 2. These regressions capture related and unrelated knowledge transfers and technology building as reflected in the propensity of licensees to cite licensor patents.[19] The results of these regressions provide broad support for our hypotheses. Looking at column 1, showing related citations, we see that the coefficient on JV is positive, consistent with the argument that a license that takes place within the context of a JV agreement leads to higher *related* technology building by the Japanese firm (H1A). The magnitude

further details. Limiting patents used in constructing the measure to those applied for during the 5 years prior to the sample period does not materially change the results.

[19] Because licensee size and profit data are available for only 521 of the 536 alliances in our sample, the regressions in Table 2 have $N = 521$. Estimations using the full sample of 536 alliances (omitting the size and profit variables) yield statistically identical results for all of the hypothesized effects.

Table 1 Summary Statistics

	Complete sample					Dyads with joint ventures				
Variable	Obs.	Mean	Std. dev.	Min	Max	Obs.	Mean	Std. dev.	Min	Max
RELATED POSTLICENSE CITES	536	40.50	165.70	0	2,128	27	197.59	370.37	0	1,497.00
UNRELATED POSTLICENSE CITES	536	76.48	269.31	0	3,125	27	190.93	362.50	0	1,465.00
POSTLICENSE BREADTH	536	0.320	0.36	0	0.998	27	0.26	0.28	0	0.73
CITATION LAG (RELATED)	278	8.75	4.35	1.0	23	18	6.86	2.78	2.0	13.50
CITATION LAG (UNRELATED)	337	8.59	3.93	1.5	22	19	6.82	2.68	2.0	13.50
JV	536	0.05	0.22	0	1	27	1	0	1	1
CROSS-LICENSE	536	0.17	0.38	0	1	27	0.22	0.42	0	1
RELATED PRELICENSE CITES	536	12.06	57.80	0	901	27	53.04	83.20	0	223
UNRELATED PRELICENSE CITES	536	21.63	77.99	0	1,072	27	97.37	228.83	0	1,072.00
PRELICENSE BREADTH	536	0.26	0.35	0.00	0.992	27	0.28	0.35	0.00	0.92
OVERLAP	536	0.41	0.29	0.00	1.0	27	0.46	0.31	0.00	0.97
PATENT IMPORTANCE	536	32.27	57.86	0.00	307.12	27	51.60	64.54	0.13	173.26
LICENSEE PATENTS	536	3.66	4.46	0.00	13.48	27	2.97	3.87	0.00	10.53
LICENSEE SIZE	521	101,452	95,142	1,900	780,000	27	95,420	82,567	2,500	243,155
LICENSEE PROFITS	521	36,470	53,027	4,677	305,863	27	38,268	59,959	824	305,863
LICENSEE JV EXPERIENCE	536	5.43	6.00	0.00	31	27	4.85	6.10	0.00	21
LICENSOR TECH CONCENTRATION	536	0.36	0.33	0.04	1.00	27	0.27	0.29	0.04	1.00
LICENSEE TECH CONCENTRATION	536	0.21	0.20	0.04	1.00	27	0.18	0.18	0.05	1.00
TEXTILES	536	0.04	0.20	0	1	27	0.04	0.19	0	1
CHEMICALS	536	0.12	0.32	0	1	27	0.22	0.42	0	1
PHARMA	536	0.04	0.20	0	1	27	0.00	0.00	0	0
MATERIALS	536	0.02	0.14	0	1	27	0.04	0.19	0	1
METALS	536	0.08	0.27	0	1	27	0.19	0.40	0	1
MACHINERY	536	0.06	0.23	0	1	27	0.04	0.19	0	1
ELEC	536	0.34	0.47	0	1	27	0.26	0.45	0	1
TELECOM	536	0.11	0.31	0	1	27	0.07	0.27	0	1
TRANSPORT	536	0.05	0.22	0	1	27	0.07	0.27	0	1
INSTRUMENTS	536	0.08	0.27	0	1	27	0.04	0.19	0	1
OTHER MFG	536	0.03	0.18	0	1	27	0.00	0.00	0	0
SERVICES	536	0.03	0.18	0	1	27	0.04	0.19	0	1
YEAR = 1988	536	0.19	0.39	0	1	27	0.22	0.42	0	1
YEAR = 1989	536	0.27	0.45	0	1	27	0.30	0.47	0	1
YEAR = 1990	536	0.27	0.44	0	1	27	0.15	0.36	0	1
YEAR = 1991	536	0.27	0.44	0	1	27	0.33	0.48	0	1

of the effect is also quite large: postestimation calculations show that, holding all other variables at their mean, the expected level of RELATED POSTLICENSE CITES for licensees in a joint venture is 140% higher than for licensees in bare license agreements. However, the estimate is also quite imprecise: in a two-tailed *t*-test, the result is significantly different from zero only at the 10% level. The coefficient on CROSS-LICENSE in this regression is also positive and slightly smaller than for JV, but the two coefficients are not statistically distinguishable. Thus, there is only partial support for H1B here: cross-license agreements result in higher transfer of alliance-related knowledge than a unilateral license, but the three alliance modes (unilateral license, cross-license, and JV) are not neatly arrayed along a continuum with respect to alliance-related knowledge transfers, as H1B implies.

Turning to the results in column 2, where the dependent variable is UNRELATED POSTLICENSE CITES, we see evidence consistent with H3, regarding the control features of the equity joint venture structure. The negative coefficient on JV in column 3 suggests that the JV structure indeed impedes knowledge transfers in technology areas that are unrelated to the alliance activities. Remarkably, even compared with a bare unilateral licensing arrangement, the transfer of unrelated knowledge appears to be reduced when licensing occurs in the context of a joint venture: estimated postlicense cites in technology classes unrelated to the licensed patents are 41.3% lower than in an equivalent bare license. Although the coefficient estimate is again quite imprecise (significant only at the 10% level), our inference is bolstered by the observation that the coefficient on CROSS-LICENSE is in the opposite direction, indicating that unrelated knowledge flows are greater in cross-licensing agreements than they are in bare unilateral agreements. Consequently, the null hypothesis that the coefficients for JV and CROSS-LICENSE are equal can be rejected

Table 2 Related and Unrelated Knowledge Building

$N = 521$	(1) RELATED POSTLICENSE CITES	(2) UNRELATED POSTLICENSE CITES	(3) POSTLICENSE BREADTH
JV	0.877†	−0.533†	−0.123**
	(0.471)	(0.319)	(0.047)
CROSS-LICENSE	0.758**	0.501†	0.011
	(0.278)	(0.284)	(0.035)
PRELICENSE CITES[a]	0.008***	0.002	0.482**
	(0.002)	(0.001)	(0.055)
OVERLAP	2.638***	3.131***	0.141***
	(0.398)	(0.344)	(0.045)
PATENT IMPORTANCE	0.015***	0.021***	0.001***
	(0.005)	(0.006)	(0.000)
LICENSEE PATENTS	0.231***	0.276***	0.009*
	(0.051)	(0.049)	(0.004)
LICENSEE SIZE	−0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
LICENSEE PROFITS	0.000	−0.000	−0.000
	(0.000)	(0.000)	(0.000)
LICENSEE JV EXPERIENCE	−0.017	−0.015	0.000
	(0.025)	(0.020)	(0.002)
LICENSOR TECH CONCENTRATION	−0.425	−2.350***	−0.217***
	(0.562)	(0.598)	(0.043)
LICENSEE TECH CONCENTRATION	−1.971**	−2.003**	−0.100†
	(0.674)	(0.641)	(0.054)
CONSTANT	1.007	2.457***	0.195**
	(0.819)	(0.645)	(0.074)
INDUSTRY DUMMIES	Included**	Included***	Included*
YEAR DUMMIES	Included**	Included**	Included
Log pseudo-likelihood	−1,385.1	−1,519.5	
Wald chi-square (25)	616.3***	826.6***	
R-square			0.600

Note. Robust standard errors are in parentheses, with clustering on licensor; number of clusters = 254.

[a]Prelicense cites measure is RELATED for column 1, UNRELATED for column 2, and PRELICENSE BREADTH for column 3.

*$p = 0.05$; **$p = 0.01$; ***$p = 0.001$; †$p = 0.10$.

with a high degree of confidence ($p = 0.009$). This evidence thus provides strong support for the argument that the enhanced control and monitoring mechanisms in the equity joint venture allow licensors to bound knowledge flows and control leakage in areas that are not directly relevant to the successful execution of alliance activities.

The results in column 3 put these different effects in perspective, reinforcing the support for H1 and H3. Focusing on POSTLICENSE BREADTH—i.e., unrelated citations as a fraction of total citations—we see the net effect of alliance structure on the shape of knowledge flows and subsequent technology building by the Japanese partner. Here the coefficient on JV is negative and highly significant (1% level in a two-tail test), in line with the observation that alliance-related knowledge flows are enhanced while unrelated flows are curtailed. For CROSS-LICENSE, the coefficient is insignificant: compared with a unilateral license, there is a symmetric increase in both related and unrelated flows, such that the relative breadth of knowledge flows is indistinguishable in the two alliance types.

The coefficients on the control variables in these analyses are also of interest and are consistent across the three specifications. As expected, there is strong serial correlation in citations: postlicense citations are significantly related to the citation level prior to the sample period. Greater technological overlap also increases postlicense citation, providing further support for the concept of partner-specific absorptive capacity (Lane and Lubatkin 1998, Mowery et al. 2002). Japanese firms also cite their U.S. partners' patents more frequently when the U.S. partner holds "important" patents, as do licensees who apply for more patents in the postlicense period. As expected, more diverse (i.e., less concentrated) licensor or licensee patent portfolios increase citations in technology areas unrelated to the licensed patents. Licensee size, profitability, and joint venture experience have no significant effect. Several of the technology sector dummy variables are also significant (results suppressed for space considerations), but there is no discernible logic to cross-sector differences in citation rates: the sectors in which licensee-licensor citation appears to be lowest in the postlicense period (controlling of course for prelicense citations, etc.) are textiles, pharmaceuticals, and electronics—quite a varied group.

Table 3 shows estimation results relevant to H2 regarding the *speed* of knowledge transfer and integration by the Japanese firm and the impact of organizational form on the average citation lag in licensees' new patent applications after initiation of the license. Consistent with our theoretical arguments we see that the citation lag for JV participants in the postlicense period is significantly lower than that for participants in a bare unilateral licensing agreement (5% significance level) after controlling for lags in the prelicense period. Interestingly, this effect is similar for both alliance-related and unrelated patents—although the magnitude of the coefficient is slightly larger for related patents, the difference between the coefficients is not statistically significant. Put beside the effects observed in the previous regressions, these results tell an interesting story: Although licensors in joint ventures put significant limitations on the extent of knowledge transfers in areas unrelated to alliance activities, where such transfers *are* permitted the knowledge is integrated into the activities of the licensee more rapidly. This latter effect is consistent with the enhanced informal learning and coordination associated with the joint venture structure.

For cross-licenses, we also see smaller citation lags in the postlicense period relative to those in a unilateral license, but this effect is significant only for

Table 3 Speed of Knowledge Integration

	CITATION LAG (RELATED) $n = 273$	CITATION LAG (UNRELATED) $n = 276$
JV	−1.589*	−1.178*
	(0.775)	(0.667)
CROSS-LICENSE	−1.887**	−0.758
	(0.700)	(0.485)
PRELICENSE CITATION LAG[a]	0.064	0.367**
	(0.157)	(0.118)
OVERLAP	−0.881	0.713
	(1.395)	(1.019)
PATENT IMPORTANCE	−0.011**	−0.009**
	(0.005)	(0.004)
LICENSEE PATENTS	0.076	−0.069
	(0.116)	(0.068)
LICENSEE SIZE	−0.000	−0.000
	(0.000)	(0.000)
LICENSEE PROFITS	0.000	−0.000
	(0.000)	(0.000)
LICENSEE JV EXPERIENCE	0.087†	0.056†
	(0.049)	(0.032)
LICENSOR TECH CONCENTRATION	−3.197†	−3.464**
	(1.818)	(1.151)
LICENSEE TECH CONCENTRATION	−2.414	0.919
	(2.407)	(2.709)
CONSTANT	9.309***	9.303***
	(2.236)	(2.002)
INDUSTRY DUMMIES	Included**	Included
YEAR DUMMIES	Included†	Included**
F	2.78***	4.47***
R-square	0.184	0.252

Note. Robust standard errors are in parentheses, with clustering on licensor.
[a]Prelicense citation lag is RELATED for column 1 and UNRELATED for column 2.
*$p = 0.05$; **$p = 0.01$; ***$p = 0.001$; †$p = 0.10$.

related citations. Indeed, postlicense citation lags for cross-licenses are indistinguishable from those in joint ventures in related areas. None of the theoretical perspectives discussed in the paper provide a complete explanation for this result: as discussed earlier, there is no reason to believe that cross-licenses are more "organizationally embedded" than unilateral licenses, and, from a transaction cost perspective, they should be at best intermediate between bare unilateral licenses and joint ventures in incentive alignment and the adoption of high-bandwidth communication channels. We return to this puzzling result in the concluding section of the paper.

5. Alternative Explanations and Limitations

Before discussing implications for research and practice, we assess alternative explanations for the observed empirical regularities, explore the robustness of our empirical results, and discuss limitations of our findings.[20]

One potential concern about our empirical analysis is the validity of patent citations as a measure of knowledge transfer. Although this use of patent citation data has a long history in the technology strategy literature (see, e.g., Almeida 1996, Mowery et al. 1996, Almeida et al. 2002, Sampat et al. 2003, Fabrizio 2007), recent research on examiner-added citations has questioned whether customary interpretations are appropriate in all cases. Alcacer and Gittelman (2006), for example, find that a significant proportion of patent citations are added by patent examiners rather than by employees of the innovating firm, suggesting that citations listed in a patent reflect not only the innovating firm's knowledge but also additional information perhaps unrelated to that knowledge base. A common assumption in prior work is that examiner-added citations simply constitute noise in an otherwise unbiased signal of knowledge flows; support for this position can be found in the results of the Jaffe et al. (2000) study showing that inventors are more likely to know about works cited in their own patents than those in a control group of "placebo" patents. Alcacer and Gittelman (2006) show, however, that examiner-added citations differ systematically from citations recorded by the innovating firm, in that examiners tend to more frequently add cross-border citations, cites in the same technology class, cites to more recent patents, and (counterintuitively) self-citations. Fortunately we have no reason to believe that these differences vary systematically across the alliances in our sample, and so they should not introduce any significant bias—rather, we conclude that examiner-added citations at worst increase the noisiness of our citation measures.

The presence of examiner-added cites is also likely to mitigate any impact of strategic citation behavior, which might otherwise be a concern. Here the specific concern is that a licensor may be less inclined to sue its joint venture partner relative to other licensees and that this may bias the observed citation patterns if citation changes the probability or expected outcome of litigation. However, as noted elsewhere (Lampe 2008), adding citations likely has two competing litigation-related effects on the value of a new patent: (i) reducing the likelihood that the new patent will be invalidated by a rival's legal challenge and (ii) increasing damages payable to the rival should the firm be found to have "willfully" infringed on a cited patent. Lampe (2008) argues that reduced litigation risk thus has an ambiguous effect on the incentives of a licensee to cite a licensor's patents; he also

[20] We thank the anonymous reviewers for several of the insights in this section.

notes that the inclusion of examiner-added cites actually undermines any attempts by the patenting firm to manipulate citations in this way. Based on this prior evidence we can therefore rule out strategic citation as a likely source of bias in our empirical results.

Beyond these general questions regarding the meaning and interpretation of patent citation rates, one might question how citation patterns emerge specifically in the context of Japanese in-licensing, particularly for joint ventures. For bare license agreements, the genesis of innovations covered in new patents applied for by the Japanese licensee after the license is established is unambiguous, and associated citations to the licensor's patents clearly indicate that the licensee's innovations build on the licensor's technology. Things are not necessarily so clear-cut in the case of an equity joint venture, however, where one could imagine that innovations created within the joint venture might result in patents that are then assigned to (i.e., owned by) either parent, or even the joint venture itself. The first thing to note in this regard is that it is very rare for joint ventures themselves to own patents. Indeed, only one JV in our sample itself held patents, and none of these patents generated any relevant citations that could contaminate our analysis. Instead, patents are typically assigned to one or other of the parent companies according to criteria stipulated in the joint venture agreement.[21] In our data, we are unable to distinguish between any patents granted to the Japanese licensee that are linked directly to JV activities and patents resulting from the firm's independent innovative activities. However, the inferences that can be drawn from our analysis do not rely on this distinction: regardless of whether an innovation is generated within the joint venture or beyond, ownership of the associated patent(s) by the Japanese firm means that they may apply the technological knowledge embodied in the innovation across the rest of their business. As such, citations to the licensor's technologies that are generated by any of these patents still accurately reflect an increase in the extent to which the Japanese licensee is building on (and can continue to build on) the licensor's technology.

Another alternative explanation for our observed results relates to the time-frame of our study (1988–1992) and the rapid increase in industrial computerization around this time. It is possible that a significant fraction of the licenses in our sample involved the acquisition of information and communications technology (ICT) for incorporation into existing products or to improve production processes. If the Japanese firms' products were themselves unrelated to ICT technologies and such licenses mainly take place outside of JV agreements, then this could possibly explain the negative coefficient on JV in our unrelated citation and breadth regressions. To investigate this possibility we returned to the data and identified ICT-related licenses. Counter to our suspicions, these licenses were just as likely to be organized as a JV in our sample as were licenses for other types of technology. Furthermore, adding a dummy variable indicating the inclusion of ICT patents in the focal license agreement does not materially change the results of our breadth regression.[22] We can therefore effectively rule out this alternative explanation.

Finally, we must address the fact that the phenomenon we are studying—like all organizational design problems—exists within a broader endogenous system of organizational and strategic choices. In our context, for example, we know that firms selectively adopt different alliance forms, in part based on concerns regarding knowledge leakage (Oxley 1997, Colombo 2003). To the extent that there are systematic but unobserved differences in the characteristics of licensees and licensors entering into joint ventures versus bare license agreements, or in the technologies involved in the agreements, this could lead to bias in our regression results (Shaver 1998, Hamilton and Nickerson 2003, Bascle 2008). In terms of observable characteristics, the descriptive statistics in Table 1 suggest that there are in fact few differences in firm size, profitability, and technological capabilities between licensor-licensee dyads with joint ventures and those without. Furthermore, our inclusion of an extensive set of control variables in the regressions allows us to account for many of the most obvious sources of licensor and licensee heterogeneity. Nonetheless, it is impossible to be exhaustive, and so some degree of unobserved heterogeneity and potential for omitted variables bias inevitably remain.

Ideally, to address this issue of endogenous selection, we would like to identify an exogenous process that influences the organizational variable of primary interest here (JV) but that does not influence knowledge flows in an alliance. This would then allow

[21] For the types of agreements in our sample—JVs created to manufacture products in Japan based primarily on U.S. technology—the U.S. partner will usually retain ownership rights to the technology brought to the joint venture. The U.S. partner will also typically retain rights to any follow-on technologies that result directly from JV operations, but there may be some cases where the Japanese firm acquires patent rights for an innovation created in the JV. This suggests that patent assignment rules may themselves represent one important knowledge control mechanism in joint ventures. Exploration of these rules and how they are determined would be an interesting avenue for further research.

[22] Details of the procedure used to isolate ICT-related patents and of the empirical analysis can be found in Table EC. 1 in the online appendix (provided in the electronic companion to this paper, which is available as part of the online version that can be found at http://mansci.journal.informs.org/).

us to instrument for the endogenous organizational choice and effectively correct for omitted variable bias. Unfortunately, finding conceptually sound *and* statistically valid instruments in this setting is practically impossible. As such, it is unsurprising that none of the previous papers on knowledge flows in alliances directly tackles the endogeneity issue, even though it is recognized as a concern and a limitation in some of the most recent work (Gomes-Casseres et al. 2006).

Although our empirical setting essentially rules out identification of conceptually motivated instruments, we searched for and found two variables that satisfy the *statistical* requirements for valid instruments for the JV variable.[23] These are licensee sales revenue and *total* prelicense cites. Estimating a two-stage model of postlicense breadth using these variables as instruments for JV allows us to assess possible bias in the results reported above. The estimation results for the two-stage regression[24] are quite consistent with the OLS results. If anything, bias in the OLS regression appears to work against us because the JV coefficient is smaller (less negative) in the OLS regression than in the two-stage least-squares regression. However, a Hausman specification test indicates that the coefficients in the two regressions are not systematically different. Thus we conclude that, although our inability to fully confront the endogeneity issue represents an important limitation of our analysis, there does not appear to be any evidence of significant bias in our results.

6. Implications and Extensions

Despite sustained interest in alliances as knowledge acquisition mechanisms from researchers in different strategy traditions, few attempts have been made to probe and refine the various perspectives on these ventures by applying them simultaneously in large-scale empirical analysis. As a result, researchers in the different traditions have tended to either talk past each other or engage in conceptual debates that are seldom if ever resolved (e.g., Kogut and Zander 1992, Foss 1996). The empirical analysis reported here builds on and extends previous research in these different traditions, adding new insights into the organizational mechanisms shaping the scope of knowledge transfer. The picture that emerges from our empirical analysis is one of particularly intense but "contained" knowledge transfer in equity joint ventures relative to bare license agreements: knowledge transfers directly related to the alliance activity are enhanced in the JV, and the speed of integration into the Japanese firm's subsequent innovations also increases. In marked contrast, leakage of technology unrelated to the alliance is significantly reduced in the JV.

Our results reaffirm the importance of incentive alignment and high-bandwidth communication channels in facilitating knowledge transfer but also highlight the importance of the joint venture's administrative structures for effectively bounding knowledge flows. This in turn is indicative of the important role played by organizational structure in mitigating potential opportunism and creating mutual commitment in alliances: the U.S. licensors in our study appear to be more open to sharing alliance-related knowledge when they have confidence that knowledge transfers are limited in unrelated domains. Thus, although the intensity of cooperation in the equity joint venture speaks to coordinative benefits associated with the shared language and routines that come with organizational embeddedness, our findings are not consistent with a "pure" knowledge-based (and opportunism-free) perspective on alliances.

We believe that our analysis also has the potential to reinvigorate interest in the unique control aspects of equity joint ventures and other alliance structures, both for their impact on the scope of knowledge transfer and for interfirm cooperation more generally. A basic assumption of our study—and one that it shares with other research on alliance governance—is that the control features of an equity joint venture cannot easily be replicated by partners in a contractual alliance. Japanese legal rules and cultural norms provide substantial support for this contention in our empirical context, but we cannot be sure that the assumption holds more generally: Despite its centrality in theories of governance, the issue has drawn little critical attention. Thus, although it is clearly beyond the scope of the current paper, more research is needed to understand the extent to which the control structures found in a joint venture can be replicated in contractual agreements and how this might vary in different institutional contexts.

Our results on the breadth and speed of knowledge transfers in cross-license agreements also raise issues worthy of further investigation. In some respects, cross-licenses appear to lie between unilateral licenses and equity joint ventures, as transaction cost economists would predict, but in other respects—and in particular in the speed of knowledge integration for alliance-related technologies—they are indistinguishable from the JVs in our sample. One might speculate that administrative controls are most important for bounding the scope of knowledge flows whereas

[23] Valid instruments are ones that are (1) correlated with the "troublesome" explanator—here JV; (2) not explanators in the original equation; and (3) uncorrelated with the error term (Murray 2006).

[24] See the online appendix, Table EC.2, for details of the two-stage least squares regression and associated specification tests (tests of instrument relevance, overidentification tests of instrument validity, and a Hausman specification test).

informal working relationships determine the intensity of interactions and the speed of knowledge transfer. Given the nature of our data we are not able to fully explore the processes by which knowledge migrates between the partner firms—and in particular how knowledge flows into and out of a joint venture are shaped and controlled. Exploration of these questions must await access to data marrying systematic information on formal organizational structures and contracts with rich detail on operational practices in alliances—a challenging requirement.

Future research may also generate more nuanced insights into the relationship between the breadth of knowledge transfers and interpartner competition. Prior research suggests, for example, that the potential consequences of technology leakage may vary depending on the competitive relationship between alliance partners (Oxley and Sampson 2004). In cases where alliance partners do not compete directly, even broad knowledge transfers may be relatively benign. Conversely, if alliance partners are very close competitors, then even quite "narrow" transfers not explicitly covered in the license may have adverse consequences for the licensor. Again this issue is beyond the scope of our study (and beyond the reach of our data), but we note it as a fascinating issue for further research.

Access to more extensive data covering a longer time-frame would also allow extensions of the current research to additional questions related to the *timing* of changes in patenting and citation behavior after alliance formation and could also help to isolate the effects of individual alliances in ongoing technology acquisition programs. Ultimately our aim must be to push toward deep understanding of how individual organizational decisions interact with the broader endogenous system of strategic choices. Although we believe that we have successfully allayed the most serious concerns about endogeneity bias in our empirical analysis, definitive treatment requires data of exceptional depth and breadth. As we look to future developments along these lines, propensity-scoring methods may prove to be a useful tool for addressing the endogeneity of strategic choice and knowledge flows, as illustrated by MacGarvie's (2006) study examining the effect of international trade on firms' technology-building activities.[25] The requirements for detailed panel data make implementation of the propensity-scoring methodology infeasible in our empirical context, but it constitutes another very useful direction for future research in this area.

We have come a long way in our understanding of international alliances since the days when commentators warned that alliances between U.S. and Japanese firms would "give away our future" (Reich and Mankin 1986). Nonetheless, as international collaboration continues apace—particularly with companies in emerging economies such as China—there is enduring interest in how firms can capitalize on alliance activity while avoiding unintended knowledge leakage. By distinguishing between knowledge flows in alliance-related technological areas and flows in unrelated domains we are able to bring greater theoretical clarity to this issue. Our study provides important new evidence that the equity joint venture structure enhances sharing of alliance-related knowledge in part *because* it can help firms to limit unintended leakage. As such, our study illuminates continuing debates about incentive alignment, organizational embeddedness, and control in alliances and provides an important foundation for future research deepening our understanding of interfirm collaboration and technology development.

[25] MacGarvie (2006) examines the effect of importing and exporting on patent citation patterns. She first estimates the effect of various firm location and resource characteristics on the propensity to import or export and then examines how the influence of foreign technology differs between pairs of firms (e.g., one importer and one nonimporter), matched by propensity scores, in a differences-in-differences specification.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

This paper is based in part on NISTEP Discussion Paper No. 34, 2004. The authors thank the National Institute of Science and Technology Policy (NISTEP) for allowing the second author access to the licensing data used in the study. The second author also acknowledges the Ministry of Education, Culture, Sports, Science and Technology (MEXT) for financial support (Subject Nos. 14730106 and 18530295). The authors thank Scott Shane, the associate editor, and two anonymous referees for their insights and guidance in developing the paper. The authors are also grateful to the many colleagues and friends who commented on earlier versions of the manuscript; all remaining errors are the authors'.

References

Ahmadjian, C. L., J. E. Oxley. 2006. Using hostages to support exchange: Dependence balancing and partial equity stakes in Japanese automotive supply relationships. *J. Law, Econom., Organ.* **22**(1) 213–233.

Alcacer, J., M. Gittelman. 2006. How do I know what you know? Patent examiners and the generation of patent citations. *Rev. Econom. Statist.* **88**(4) 774–779.

Allen, S. A. 1991. Caterpillar-Mitsubishi (A). Babson College Case Study 392-090-1, Babson College, Wellesley, MA.

Almeida, P. 1996. Knowledge sourcing by foreign multinationals: Patent citation analysis in the U.S. semiconductor industry. *Strategic Management J.* **17**(Winter) 155–165.

Almeida, P., J. Song, R. Grant. 2002. Are firms superior to alliances and markets? An empirical test of cross border knowledge building. *Organ. Sci.* **13**(2) 147–161.

Balakrishnan, S., M. P. Koza. 1993. Information asymmetry, adverse selection and joint ventures. *J. Econom. Behav. Organ.* **20** 99–117.

Bascle, G. 2008. Controlling for endogeneity with instrumental variables in strategic management research. *Strategic Organ.* **6**(3) 285–328.

Colombo, M. G. 2003. Alliance form: A test of the contractual and competence perspectives. *Strategic Management J.* **24**(12) 1209–1229.

Dernis, H., M. Khan. 2004. Triadic patent family methodology. OECD STI Working Paper 2004(2), Organisation for Economic Co-operation and Development Science, Technology, and Industry, Paris.

Elkjaer, B., 2003. Social learning theory: Learning as participation in social processes. M. Easterby-Smith, M. A. Lyles, eds. *The Blackwell Handbook of Organizational Learning and Knowledge Management*. Blackwell Publishers, Oxford, 38–52.

Fabrizio, K. 2007. University patenting and the pace of industrial innovation. *Indust. Corporate Change* **16**(4) 505–534.

Foss, N. J. 1996. More critical comments on the knowledge-based theories of the firm. *Organ. Sci.* **7**(5) 519–523.

Geringer, J. M., L. Hebert. 1989. Control and performance of international joint ventures. *J. Internat. Bus. Stud.* **20**(2) 235–254.

Gilson, R. J., M. J. Roe. 1993. Understanding the Japanese kieretsu: Overlaps between corporate governance and industrial organization. *Yale Law J.* **102** 871–906.

Gomes-Casseres, B. 1991. Xerox and Fuji Xerox. Harvard Business School Case Study 9-391-156, Harvard Business School Press, Boston.

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

Hall, B. H., A. B. Jaffe, M. Trajtenberg. 2001. The NBER patent citation data file: Lessons, insights and methodological tools. NBER Working Paper 8498, National Bureau of Economic Research, Cambridge, MA.

Hamilton, B., J. A. Nickerson. 2003. Correcting for endogeneity in strategic management research. *Strategic Organ.* **1** 53–80.

Heiman, B., J. A. Nickerson. 2004. How do firms manage knowledge sharing while avoiding knowledge expropriation in interfirm collaborations? *Managerial Decision Econom.* **25** 401–420.

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

Jaffe, A. B., M. Trajtenberg. 2002. *Patents, Citations and Innovations*: *A Window on the Knowledge Economy*. MIT Press, Cambridge, MA.

Jaffe, A. B., M. Trajtenberg, M. S. Fogarty. 2000. Knowledge spillovers and patent citations: Evidence from a survey of inventors. *Amer. Econom. Rev.* **90**(2) 215–218.

Janowicz-Panjaitan, M., N. Noorderhaven. 2008. Formal and informal interorganizational learning within strategic alliances. *Res. Policy* **37** 1337–1355.

Kogut, B. 1988. Joint ventures: Theoretical and empirical perspectives. *Strategic Management J.* **9** 319–332.

Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities and the replication of technology. *Organ. Sci.* **3**(3) 283–301.

Lampe, R. 2008. Strategic citation. Working paper, Stanford University, Stanford, CA.

Lane, P. J., M. Lubatkin. 1998. Relative absorptive capacity and interorganizational learning. *Strategic Management J.* **19**(5) 461–477.

Liebeskind, J. P. 1996. Knowledge, strategy, and the theory of the firm. *Strategic Management J.* **17**(Winter) 93–107.

MacGarvie, M. 2006. Do firms learn from international trade? *Rev. Econom. Statist.* **88**(1) 46–60.

Mowery, D. C., J. E. Oxley, B. S. Silverman. 1996. Strategic alliances and inter-firm knowledge transfer. *Strategic Management J.* **17**(Winter) 77–91.

Mowery, D. C., J. E. Oxley, B. S. Silverman. 2002. The two faces of partner-specific absorptive capacity: Learning and co-specialization in strategic alliances. F. Contractor, P. Lorange, eds. *Cooperative Strategies and Alliances*. Elsevier, London, 291–319.

Murray, M. P. 2006. Avoiding invalid instruments and coping with weak instruments. *J. Econom. Perspectives* **20**(4) 111–132.

Osborn, R. N., C. C. Baughn. 1990. Forms of interorganizational governance for multinational alliances. *Acad. Management J.* **33**(3) 503–519.

Oxley, J. E. 1997. Appropriability hazards and governance in strategic alliances: A transaction cost approach. *J. Law, Econom., Organ.* **13**(2) 387–409.

Oxley, J. E. 2004. Learning versus protection in interfirm alliances: A false dichotomy. V. Mahnke, T. Pedersen, eds. *Knowledge Flows, Governance and the Multinational Enterprise*. Palgrave Macmillan, New York, 108–129.

Oxley, J. E., R. C. Sampson. 2004. The scope and governance of international R&D alliances. *Strategic Management J.* **25** 723–749.

Pisano, G. 1989. Using equity participation to support exchange: Evidence from the biotechnology industry. *J. Law, Econom., Organ.* **5**(1) 109–126.

Reich, R. B., E. D. Mankin. 1986. Joint ventures with Japan give away our future. *Harvard Bus. Rev.* **64**(2) 78–86.

Reuer, J., A. Arino, T. Mellewigt. 2006. Entrepreneurial alliances as contractual forms. *J. Bus. Venturing* **21**(3) 306–325.

Sampat, B., D. C. Mowery, A. A. Ziedonis. 2003. Changes in university patent quality after the Bayh-Dole act: A re-examination. *Internat. J. Indust. Organ.* **21** 1371–1390.

Santoro, M. D., J. P. McGill. 2005. The effect of uncertainty and asset co-specialization on governance in biotechnology alliances. *Strategic Management J.* **26** 1261–1269.

Shaver, J. M. 1998. Accounting for endogeneity when assessing strategy performance: Does entry mode choice affect FDI survival? *Management Sci.* **44**(4) 571–585.

Tallman, S. 2003. The significance of Bruce Kogut's and Udo Zander's article, 'Knowledge of the firm and the evolutionary theory of the multinational corporation'. *J. Internat. Bus. Stud.* **34**(6) 495–497.

Yan, A., B. Gray. 1994. Bargaining power, management control and performance in United States-China joint ventures: A comparative study. *Acad. Management J.* **37** 1478–1994.

Zander, U., B. Kogut. 1995. Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test. *Organ. Sci.* **6**(1) 76–92.