Strategic Management Journal
Strat. Mgmt. J., **31**: 1183–1201 (2010)
Published online EarlyView in Wiley Online Library (wileyonlinelibrary.com) DOI: 10.1002/smj.852
Received 19 June 2008; *Final revision received 12 February 2010*

THE HIDDEN COSTS OF OUTSOURCING: EVIDENCE FROM PATENT DATA

MARKUS REITZIG[1]* and STEFAN WAGNER[2]
[1] *London Business School, London, U.K.*
[2] *INNO-tec, University of Munich, Munich, Germany*

Drawing on patent data for approximately 500 firms over 20 years, we advance recent theory on firm boundaries and test these propositions for the first time. We first provide evidence for the existence of knowledge complementarities between vertically related activities in a firm's value chain by showing that firms face increasing (decreasing) performance in conducting downstream activities (i.e., patent litigation) the less (more) they outsource related upstream activities (i.e., patent filing). We then propose and empirically demonstrate that vertical integration benefits through learning differ from vertical outsourcing costs through forgetting. We show that firms can partly offset these hidden outsourcing costs by sourcing similar upstream products from internal and external suppliers. Copyright © 2010 John Wiley & Sons, Ltd.

INTRODUCTION

A fundamental question in strategy research is how firms should draw their boundaries. One major problem with the dominating, transaction cost economics (TCE) viewpoint is that it focuses on one transaction at a time and therefore lacks a systemic approach (Argyres and Liebeskind, 1999; Puranam, Gulati, and Bhattacharya, 2008). The knowledge-based view (KBV) of the firm addresses this shortcoming. One of its key contributions is the distinction between an activity A and the knowledge required to profit from that same activity, say, K_A. By focusing on knowledge bases rather than on isolated transactions, the KBV offers a more realistic way to analyze firm boundaries. Among the important contributions drawing on the KBV that refine the simple TCE-based pictures of firm boundaries (e.g., Dyer and Hatch, 2006; Dyer and Singh, 1998; Hatch and Dyer, 2004; Helper, MacDuffie, and Sabel, 2000; Parmigiani, 2007; Rothaermel, Hitt, and Jobe, 2006; Santos and Eisenhardt, 2005), Brusoni, Prencipe, and Pavitt (2001) and Jacobides and Billinger (2006) are most directly related to this research. Both provide instances of knowledge-based complementarities between internal and external sourcing (Puranam *et al.*, 2008) that unfold across *different* layers of a firm's value chain. Brusoni *et al.* (2001) provide one important learning-based argument for why corporations do not outsource their entire production of a given upstream activity, call it A, to more efficient markets despite the lack of transactional hazards. The authors find that, in times of technological change, excessive outsourcing of supply technologies creates opportunity costs of not learning about alterations in the supply segment through internal production. These (opportunity) costs of outsourcing upstream materialize further downstream in the firm's value chain and are reflected in the firm's lessened ability to integrate external supplies into its core downstream

Keywords: knowledge-based view; patents; outsourcing
* Correspondence to: Markus Reitzig, London Business School, Sussex Place, Regent's Park, London, NW1 4SA, UK.
E-mail: mreitzig@london.edu

production activities, say B. Hence, failing to do A leads to a decrease in the knowledge required to perform well in doing B ($\downarrow A \Rightarrow \downarrow K_B$), which we denote as 'vertical outsourcing costs' in this paper. Jacobides and Billinger (2006) find support for the same type of knowledge-based complementarity between activities that are vertically related in the value chain—however, within a different empirical scenario. The authors show that partially integrating a novel upstream supply activity A may help a firm adjust its downstream production knowledge base K_B, thereby increasing the marginal value of carrying out production activities B in-house. The authors coin the term 'dynamic benefits of vertical architecture' (hereafter also called: dynamic benefits of vertical integration) in order to describe how the knowledge complementarities between doing A and benefitting from B unfold over time ($\uparrow A_t \Rightarrow \uparrow K_{B,t+\tau}$).

These intriguing insights about the existence of knowledge complementarities across value chain activities trigger further research questions. First, along a basic *empirical* dimension, large-scale evidence for the existence of cross-activity, knowledge-mediated complementarities is currently lacking. As scholars we would like to see the fundamental vertical complementarity argument underlying earlier case-based works to withstand a more robust econometric test. Second, and more importantly from a *theoretical* standpoint, the question emerges whether vertical outsourcing costs (Brusoni *et al*., 2001) and dynamic benefits of vertical integration (Jacobides and Billinger, 2006) are always flipsides of the same coin as standard neoclassical economics would suggest. In fact, the prior study of both vertical outsourcing costs and vertical integration benefits was conducted in settings whereby firms would benefit from the acquisition of new knowledge in one way or another. However, such scenarios in which the firm's *learning* (Cohen and Levinthal, 1990; Argote, 1996) is the mechanism by which knowledge complementarities arise may differ substantially from those in which firms are concerned with actually *forgetting* (Walsh and Ungson, 1991; Olivera, 2000) existing knowledge in the process of outsourcing. At present, however, limited knowledge exists about such hidden vertical outsourcing costs through forgetting (henceforth also called 'hidden outsourcing costs'), whether they exist, and how they may differ from learning-based benefits of vertical integration.

In this paper we attempt to address the two scientific gaps previously described by (1) providing a large-scale study of knowledge complementarities in general, and (2) analyzing in more detail the hidden outsourcing costs through forgetting. We study knowledge complementarities between outsourcing upstream and downstream activities within a firm's value chain on an unbalanced panel dataset of approximately 500 firms over 20 years, comprising more than 6,000 firm-year observations. We embed our empirical tests in the firm's intellectual property (IP) value chain (Reitzig and Puranam, 2009; Somaya, Williamson, and Zhang, 2007). Drawing on European patent examination and litigation data for European firms from the period 1980 through 2001, we analyze how the degree to which a firm outsources patent filing activities early in its value chain reduces its (related) ability to detect IP competitors at a later stage of its value chain. In order to deliver on our first research goal, we argue and empirically demonstrate that the firm's later ability to detect IP competitors is negatively affected by the loss of relevant technological and legal knowledge that results from outsourcing preceding patent filing activities across firms. In order to deliver on our second research goal, we theoretically propose that vertical outsourcing costs through forgetting, if they exist, must show a distinct feature that differentiates them from vertical integration benefits through learning. Namely, variation in internal and external knowledge bases should strictly add to hidden outsourcing costs, thereby differentiating forgetting from learning (Schilling *et al*., 2003; Sorenson and Sorensen, 2001). We provide empirical evidence for this theoretical conjecture and by doing so we manage to both show that hidden forgetting-related costs prevail in our data and characterize them.

The rest of the paper follows a standard structure. We present our theoretical arguments and hypotheses, followed by a description of our data, empirical results, and discussion.

THEORY AND HYPOTHESES

Firm scope, knowledge complementarities, and learning from doing

The understanding of the vertical scope of the firm has been shaped by different leading theories. Classic TCE (Coase, 1937; Williamson, 1971, 1985)

DOI: 10.1002/smj

frames firm boundary decisions as 'make *or* buy' problems. At the level of the individual transaction, the advantages firms gain by using external markets and benefiting from suppliers' economies of scale and specialization must be traded against the transactional risks such as holdup. Notwithstanding its explanatory power, TCE shows systematic limitations in that it neglects considerations that involve knowledge-related interdependencies between different transactions. The KBV (Grant, 1996) complements TCE in that it offers a more systemic gateway to understanding the vertical scope of the firm (Argyres and Liebeskind, 1999). Importantly, by focusing on the knowledge required to benefit from a transaction rather than on the transaction itself, the KBV helps to explain why firms engage in activities that appear inefficient from a TCE perspective. For example, a firm may source products and services in-house for which a hazard-free external market exists, if engaging in this in-house production creates knowledge that benefits the firm (Parmigiani, 2007). Where carrying out a given activity A increases the marginal value of a related activity B that also affects the firm's performance, scholars speak of knowledge complementarities.

Knowledge complementarities can unfold in various ways. For the purpose of this paper a distinction into two categories appears helpful. On the one hand, the marginal effect of conducting activity A on the value of carrying out activity B may unfold within the same layer of a firm's value chain. As an example, internal and external suppliers may share knowledge of improvements in internal (A) and external (B) production processes and technologies, thus enabling each other to enhance their efficacy and effectiveness (Dyer and Hatch, 2006; Dyer and Singh, 1998; Gulati, Lawrence, and Puranam, 2005; Hatch and Dyer, 2004; Helper *et al.*, 2000). For the sake of brevity, we denote such relationships as 'horizontal complementarities' throughout this paper.

On the other hand, recent theory has also conjectured that knowledge complementarities can unfold across different layers of the value chain—these are henceforth called 'vertical complementarities.' Here, conducting a given upstream activity increases the marginal value of engaging in a different downstream activity in a firm's value chain—or *vice versa*. Case-based evidence for this type of complementarity has been provided for two scenarios that appear to be mirror images of one another. On the one hand, firms' performance of conducting their core downstream activity has been shown to decrease when they excessively outsource knowledge-coupled upstream activities (Brusoni *et al.*, 2001) and to thus create outsourcing costs to the firm, which we call vertical outsourcing costs. On the other hand, firms may become better at carrying out a focal downstream activity when partially integrating a related upstream activity. The latter benefits are referred to in the literature as 'dynamic benefits of vertical architecture' (Jacobides and Billinger, 2006).

In this paper we focus on vertical knowledge complementarities. With regard to these complementarities, current theory predicts a negative relationship between the rate of outsourcing upstream activities and downstream performance, all else being equal. We will look at this relationship as a linear one, bearing in mind, however, that the existence of more complex functional forms must not be dismissed offhand.[1] We thus posit in sufficiently broad terms:

> *Hypothesis 1: All else being equal, increasing (decreasing) the rate of outsourcing upstream activities will decrease (increase) a firm's downstream performance*

Hidden knowledge losses—the vertical costs of outsourcing through forgetting

So far, scholars have assumed the mechanism underlying vertical knowledge complementarities to be a 'learning mechanism' in the spirit of Cohen and Levinthal's (1990) seminal paper. The logic is easily understood in the case of the benefits resulting from vertical integration (Jacobides and Billinger, 2006); however, it is also central to Brusoni *et al.*'s (2001) argument. In their study of the aircraft manufacturing industry, the downstream costs incurred through excessive outsourcing of upstream control systems supplies are the opportunity costs of missed learning in times of technological change: in this case, the missed opportunity

[1] Brusoni *et al.* (2001) suggest that there may be a threshold level of excessive outsourcing beyond which downstream performance dips become particularly pronounced. Parmigiani's (2007) work makes it appear possible that integration, hybrid integration, and complete outsourcing are not ordered but discrete choices, and that different expectations on the impact of downstream performance may result altogether. We explore these options empirically without hypothesizing them explicitly, and mention them briefly in our Discussion section.

to learn about what it takes to optimally integrate digital electronics control systems into planes, as opposed to hydromechanical control devices.

However, whereas learning may be an important mechanism in understanding the *creation* of knowledge complementarities (Darr, Argote, and Epple, 1995; Kogut and Zander, 1992; Pisano, 1994), it appears equally important to consider the mechanism of forgetting when seeking to understand the consequences of *losing an existing* knowledge stock. The loss of 'organizational memory' (Cyert and March, 1963; Huber, 1991; Levitt and March, 1988; March and Simon, 1958) may severely set back the performance of firms (Parise, Cross, and Davenport, 2006).

To the extent that creating vertical knowledge complementarities through learning and losing them through forgetting are different processes; however, it may be that vertical outsourcing costs and vertical integration are no longer just mirror images of the same phenomenon, as current theory suggests. Instead, they may reflect inherent differences in their underlying mechanisms. Interestingly this theoretical possibility—that firms lose already-existing knowledge across adjacent activities as a consequence of forgetting through outsourcing—has not thus far been examined in the literature on firm boundaries, and its specificities are not well understood.

The literature that can shed light on these characteristics of organizational forgetting across value chain activities is that on organizational memory.[2] Within this domain, one stream of research appears particularly relevant when carving out the differences between learning (through integration) and forgetting (through outsourcing): the literature on knowledge-retention means. A series of knowledge-retention mechanisms have been discussed in the past (Olivera, 2000; Walsh and Ungson, 1991): routines and production rules (Cyert and March, 1963; Nelson and Winter, 1982), files (Campbell-Kelly, 1996), computer-based information systems (Huber, 1991), and individuals (Simon, 1991) and—most directly related to this paper—the products they manufacture (Hargadon and Sutton, 1997; Olivera and Argote, 1999). However, if internal production is to help prevent the forgetting of *existing* knowledge in the presence of outsourcing, the knowledge created from the remaining internal production must correspond to the knowledge potentially lost through outsourcing. In a multiproduct firm, in which internal and external sourcing happens across a range of different product categories, this insight results[3] in a prediction that holds if vertical outsourcing costs are driven by a process of forgetting: the more similar the profiles of volumes a firm outsources and retains are across product categories, the smaller the total vertically related knowledge loss will be to the firm, all else being equal (notably the overall outsourcing volume across all products).

The logic pertaining to this proposition directly follows from its four underlying assumptions, which we deem unproblematic in the context of this paper.[4] First, there is a minimal level of internal production needed to prevent forgetting in one given product category, say P_1, that increases strictly with the volume of work being outsourced in that same category.[5] Second, the functional form of this relationship between the outsourcing rate and the degree of forgetting is identical across all product categories, $P_1 - P_n$, that a firm is engaged in. Third, pending better knowledge, we assume that the magnitude of knowledge complementarities between internal and external sourcing is constant across product categories. Forth, and again pending better knowledge, we assume that different product categories draw on different knowledge bases.[6] Under these conditions, if forgetting can be avoided by internal production, and if the

[2] This partly encompasses the literature on individual memory/forgetting.

[3] The respective assumptions needed are few and realistic. We elaborate on them further below.

[4] Note that if the assumptions were not fulfilled on a given dataset, this should only translate into a conservative bias in empirical tests pertaining to Hypothesis 2, thus making it more difficult to find evidence for Hypothesis 2. We elaborate on this in our Discussion section.

[5] This assumption elaborates on earlier work (Brusoni *et al.*, 2001) suggesting that threshold levels of internal sourcing exist that allow firms to maintain an updated knowledge base that prevents them from missing learning about emerging opportunities in their market environment. We analogously argue that such thresholds of sufficient internal sourcing to prevent forgetting prior knowledge exist. Pending better knowledge, we furthermore assume that these thresholds are proportional to the amount of knowledge that can be lost through external sourcing.

[6] That is, that internal production in category P_1 can not counter knowledge losses that occur through external sourcing in category P_2.

DOI: 10.1002/smj

volume of internal production necessary to prevent forgetting is proportional to the volume outsourced for any product category (Assumption 1), it follows naturally that, in a multiproduct firm, forgetting is minimized when the volume distribution of activity across product categories is similar for internal and external sourcing, all else being equal (and provided Assumptions 2–4 also hold).

If we then define the distribution of a firm's internal sourcing across product categories upstream as its internal upstream knowledge base, and the distribution of its external sourcing across product categories upstream as its external upstream knowledge base, then similarity[7] between these two knowledge bases should help the firm counter forgetting, and consequently reduce total downstream knowledge losses. We thus posit that the following holds whenever firms run the risk of losing knowledge by increasing their overall outsourcing rate:

> *Hypothesis 2: All else being equal, the more similar the knowledge bases underlying upstream activities performed internally and externally, the higher a firm's downstream performance will be.*

Note that this logic holds *if and only if* a mechanism of forgetting drives the vertical knowledge losses. The rationale depicted in Hypothesis 2 is plausible only when the point of reference for the knowledge loss is the firm's *own prior* knowledge—as is the case when we speak of *forgetting*. As soon as one focuses on missed learning opportunities instead of forgetting, this point of reference changes to what the firm *could have learned* about changes in its technological environment had it sourced products differently (both internally and externally). In this latter case, similar shares of internal/external sourcing across the firm's product categories should not reduce knowledge losses from missed learning under a reasonable set of assumptions.[8] In turn, however, this means that if upstream similarity increases downstream performance, at least some of the vertical outsourcing costs must be driven by a mechanism of forgetting.

EMPIRICAL SETTING, DATA, AND ANALYTICAL METHODS

Analyzing corporate patenting activities—linking hypotheses to our empirical stage

In order to address our research questions, we set up our analysis in the firm's IP value chain (Reitzig and Puranam, 2009). Studying the patent-related activities of firms provides a formidable testing ground for our questions because it allows us to track knowledge-coupled vertically related activities over time (patent filing and patent enforcement) and their degree of outsourcing (degree of patents filed by external lawyers), as well as related performance measures for both activities (rate of granted patent applications and, particularly important for this paper, the rate of detected competitors).[9]

Figure 1 (upper part) depicts a firm's stylized IP value chain (adapted and refined from Reitzig, 2007, for the purpose of this paper). The IP-generation process (mainly research and development [R&D]) precedes the patent protection stage and subsequent exploitation (including branding, licensing, etc.). Patent enforcement (including reactive and proactive litigation), although still a legal activity, succeeds patent filing in the value chain. The different types of activity are highly specialized and are therefore carried out by separate individuals in many firms. However,

---

[7] Note that by definition the value of similarity is bound between 0 (when the firm does not internally source in the same product category as externally, and *vice versa*) and 1 (the ratio of internally sourced goods to externally sourced goods is the same for each product category).

[8] The extent of learning opportunities should be related to the pace of (knowledge) changes that are likely to vary across different product areas. At least this is what the entirety of the literature on life cycle differences across different industries suggests. Given the literature in the field of learning and forgetting it seems rather implausible that *new-learning* and *not-forgetting* through internal production in a given product category are perfectly symmetrical processes. Optimal new-learning should therefore not require similar distributions between external and internal sourcing across product categories but can vary depending on the pace of change in those areas. Therefore dissimilar distributions between external and internal sourcing across product categories do not necessarily lead to an increase in missed learning opportunities. Based on this reasoning, it appears that there is *no simple* functional relationship between upstream (dis)similarity and (missed) learning. Upstream (dis)similarity, the way it is defined here, is hence characteristic for forgetting prior knowledge, not for acquiring new knowledge or missing opportunities thereof.

[9] In-depth interviews we conducted with several patent experts confirm the suitability of our research design given our research objectives. More information on these exploratory interviews is available from the authors.

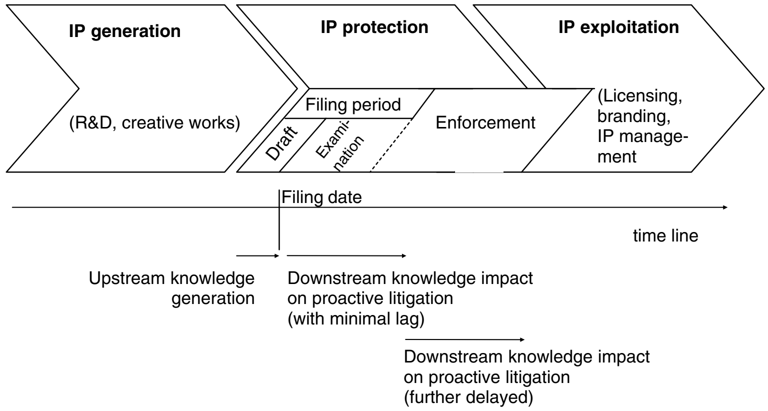

Figure 1. Firm's stylized internal intellectual property value chain with timing sequence of knowledge generation and losses

recent research also indicates that cross-functional involvement enhances the likelihood of leveraging IP successfully, suggesting that the different units interconnect by drawing on partly overlapping knowledge bases (Reitzig and Puranam, 2009; Somaya *et al*., 2007).

The link between Hypothesis 1 and our chosen setting

The vertical complementarities we study in this paper occur between (upstream) patent filing and proactive (downstream) patent enforcement activities. Patent filing and patent enforcement are two separate activities (see Figure 1, top half); however, the two draw on partly overlapping knowledge sets. More specifically, the process of drafting a patent application and filing it at the patent office entails identifying so-called prior art—generating knowledge highly relevant for the subsequent patent enforcement process. A patent's prior art encompasses all knowledge—whether it exists in written form, as a product, or contained in some other fashion—disclosed to the public before the patent was applied for (Knight, 1996; O'Connell, 2008). It constitutes the benchmark of patentability of the patent application under investigation, and allows the patent office to determine whether the current application is novel and nonobvious enough to merit patent protection.

Thus, the more (less) a firm outsources its patent filing activities, the more (less) it leaves prior art identification and interpretation to external lawyers. Prior art identification potentially carries an important interim by-product, however, which is knowledge about technology competitors. This interim information unfolds its value for the firm only when combined with further in-house knowledge of the firm's broader technology strategy. The information is of no additional value when held by external lawyers, as these legal suppliers lack the broader picture of the firm's technology strategy.[10] Information on technology competitors gained from the patent filing process crucially complements the firm's further technology intelligence,[11] enabling it to identify potential technology competitors early on and to attack them proactively as part of its enforcement strategy. Likely the most important proactive legal weapon in the European Patent Office (EPO) system is the so-called opposition procedure (Harhoff and Reitzig, 2004), which firms can use to seek invalidation of their competitors' patents at comparatively low cost. Successful oppositions require both (1) the early identification of target competitor patents (within nine months of the competitor patent being granted) and (2) an in-depth understanding of the existing prior art in the field in order to identify potential points of attack. According to the interviews we conducted for this study (more information available from the

[10] For the same reason, buying-in competitor technology intelligence from external law firms as a substitute product to compensate entirely for internal knowledge losses (Hargadon and Sutton, 1997, Olivera and Argote, 1999) proves difficult. Moreover, coordination between in-house strategists and external lawyers becomes excessively costly.

[11] Note that our estimations confirm this allegation. Even when controlling for outsourcing of downstream (litigation) activities, the effect of upstream (patent filing) outsourcing on downstream performance remains visible. Thus the detection of competitors cannot be left to external lawyers only. We return to this point in our Discussion section.

DOI: 10.1002/smj

authors), firms appear to be better on average, all else being equal, in meeting both aforementioned conditions the more patent filing they carry out in-house. Conversely, they struggle to meet either one of the conditions when excessively outsourcing their patent filing activities. In times of technological change, the performance loss may be a consequence of not staying up to date with technology knowledge. Nevertheless, even in less technologically turbulent environments, the dismissal of in-house attorneys—individual repositories of corporate memory (Simon, 1991)—leads to significant knowledge losses. Consequently, we suggest that a firm incurs performance gains (losses) related to patent enforcement the more it integrates (outsources) its patent filing services. More specifically, the number of detected potential competitors should decrease with the number of patents the firm applies for externally, all else being equal.

The link between Hypothesis 2 and our chosen setting

A crucial construct in Hypothesis 2 is knowledge base similarity between internally and externally sourced upstream activities. In our empirical setting, this corresponds to the similarity between the knowledge bases required to generate internal as opposed to external patent filings. Knowledge bases required to file patents are multidimensional. By design, the *procedural* knowledge bases (Stinchcombe, 1990) for internally and externally sourced patent filings are identical—in both cases a patent agent searches prior art and drafts a patent application, a standardized legal document. Yet the specific content of the applications, that is, the description of existing and novel technology, often differs between applications drafted in-house and externally. This is because a large fraction of patenting-active firms today are multitechnology firms. When these firms outsource a part of their patent filing activities to external law firms, different patterns of transaction heterogeneity may result depending on the firm's decision making. Take the example of any given firm that is equally active across two different technology areas TA_1 and TA_2, where TA_1 and TA_2 are the empirical representations of our more abstractly formulated product categories P_1 and P_2 (see Hypothesis 2) above. When the firm decides to outsource 50 percent of its total patent filing activities, it has a continuous spectrum of outsourcing combinations to choose from; however, the two extreme cases are the following: either the firm entirely drops its in-house patent filing activities for one of its technology areas, say TA_2, or it cuts the number of in-house applications for TA_1 and TA_2 equally across both areas and ends up outsourcing 50 percent of its filing activities to an external law firm for each of the technology areas. In the first case, the firm's internal and external transactions become more heterogeneous (He and Nickerson, 2006) along a technological dimension than in the second case, and the firm loses more of its knowledge base in one of the technology areas. In the first case, the distribution of effort spent on TA_1 and TA_2 internally will not resemble the distribution of effort spent on TA_1 and TA_2 externally, leading to dissimilar in-house and external knowledge bases. In the second case, the knowledge bases will be similar. In our data, we can distinguish between 30 different technology areas[12] that have been clustered in such a way that it appears reasonable to assume that two patent applications within one technology area will draw on far more similar knowledge than two patent applications that fall into two different technology areas.[13]

Furthermore, testing the mechanism underlying Hypothesis 2 requires an empirical stage in which hidden outsourcing costs are likely to exist. Our setting should warrant this condition, as Figure 2 illustrates. The overall outsourcing rate is increasing and hidden outsourcing costs, if they exist at all, should consequently be detectable in our data.

Finally, embedding our tests in the firm's patent-related value chain is particularly interesting for another institutional reason: there appear to be hardly any direct costs associated with outsourcing patent filing activities. Patent filing creates primarily 'efficiency' benefits for the firm (Santos and Eisenhardt, 2005), at least when one focuses on the success of patent filing as an upstream activity in isolation. As we show in a different paper (Reitzig and Wagner, 2009), a firm's sheer performance

[12] The distinction of 30 different technological areas in the EPO system dates back to the Organisation for Economic Co-operation and Development (OECD, 1994), and the categorization and its predecessors are widely used in similar studies. We use the updated version of 11 October 2000 (personal communication by Ulrich Schmoch, ISI Institute).

[13] See Hinze, Reiss, and Schmoch (1997). The cluster representation suggests that patent applications rarely draw on knowledge from two different technological areas. Some minor overlaps exist, however. In turn, this also means that our Assumption 4 should be relatively well fulfilled.

DOI: 10.1002/smj

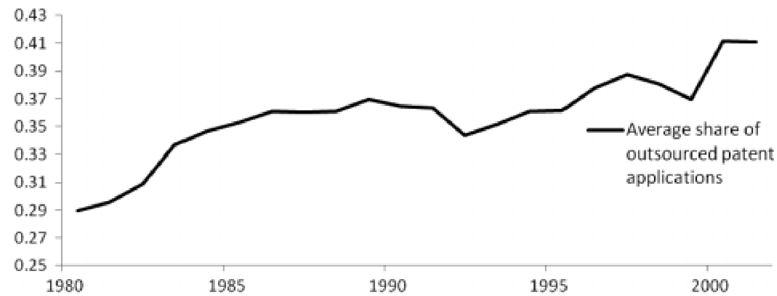

Figure 2. Outsourcing trend across the entire sample

in turning patent applications into patent grants increases linearly with the rate of outsourcing of patent applications to external law firms. Major specialization advantages of the external lawyers are likely to account for the finding. Such a setting in which no horizontal complementarities from filing patents externally and internally exist appears particularly appropriate for our analysis of vertical outsourcing costs.

Patent data

We first obtained information on individual patent application filings and then aggregated the data to the firm-year level, as we are interested in the ability of firms to successfully source IP activities. This microlevel approach not only allows us to determine the extent to which single firms engage in outsourcing activities but also yields a variety of measures related to the technical and legal characteristics of their patent applications. It is important to capture such aspects because the success of turning applications into valid patent rights depends on both the novelty of the underlying invention and the legal sophistication used to carry out the filing. The data source we used for our analyses is available through the EPO's official online source, EPOLINE Register Plus (see www.epoline.org; data extracted April 2003). We collected patent information for the 504 most active European patent applicants at the EPO.[14] Non-European applicants were excluded from our study because they are subject to EPO regulations requiring representation at the EPO by a European patent attorney (see EPO, 2000 for more details); their inclusion would potentially distort our outsourcing measure in that these applicants have no choice between integrating and outsourcing their patent-related activities.

In total, almost 1.2 million patent applications were filed at the EPO between 1980 and 2001, and European applicants account for 575,945 of these applications. The 504 firms in our sample account for about 45 percent (257,567) of the latter applications. Note that we further excluded all patent applications for which the EPO had not published a grant decision by April 2003, which left us with a final sample of 189,332 patent applications.[15] In the following section, we describe the variables we derived from this data source.

[14] Patent applicants had to have filed at least 100 patent applications at the EPO during our observational period to be included in our sample. Although this threshold may seem arbitrary, the results of our analysis are robust to different thresholds. Moreover, we put considerable effort into the cleaning of applicant names (note that the EPO may refer to the same applicant using different spellings, etc.). For all firms in our sample we consolidated names (standardizing information on legal structure and geography, and so on, to the extent that it would make sense for this analysis) using automated routines. Furthermore, a significant effort also went into the manual consolidation of subsidiary companies with parent companies for the largest 104 applicants. For these latter applicants, additional manual cross-checks on name spellings were carried out. That said, however, it is likely that our data will remain afflicted with some noise in this regard (especially for smaller firms) given its longitudinal and cross-sectional dimensions. Remaining errors should be unsystematic, however.

[15] Because the data collection was carried out in 2003, granting decisions for patents filed in the years shortly before 2001 (especially between 1998 and 2001) suffer from truncation. By dropping the pending cases, we incur a systematic selection bias that artificially increases opposition-to-filing rates in recent years. In order to test whether the bias has a major impact on the robustness of our findings, we carried out separate regressions (not reported in this paper) on subsamples of older patents for which granting decisions would be more completely observable. The findings confirm that the selection bias is of no major relevance for this paper; however, the significance of the similarity measure in Model 3B becomes less stable when turning to older patent applications.

DOI: 10.1002/smj

Variable definitions and measures

Dependent variable

We are interested in relating a firm's *downstream performance* at the firm-year level to the rate of outsourcing upstream activities. Both our interviews and prior literature (Reitzig, 2007) suggest that within the IP value chain, this downstream performance is reflected in the firm's ability to detect targets for proactive litigation.[16] We thus conceive of downstream performance as the firm's ability to identify competitor patents and litigate against third parties' IP rights, relative to the firm's effort invested into IP. Dividing the number of oppositions filed in year t (*downstream performance numerator*) by the number of applications of the focal firm filed in the same year (*downstream performance denominator*)—in order to proxy for the overall importance the firm attributes to using patents—appears to be a suitable measure. For methodological reasons, which we explain in more detail below, our actual dependent variable becomes the firm-year *downstream performance numerator*, a count variable (see 'Model Specifications and Econometric Issues'). The firm-year *downstream performance denominator* enters as an independent variable in our regressions.

Independent variables

We are interested in measuring the *outsourcing rate of upstream activities* that increase the value of engaging in a downstream activity. Within the IP value chain, engaging in upstream patent filing increases the firm's ability to identify competitor patents (see prior discussion). We thus compute the firm-year rate of patent filing outsourcing as our relevant upstream outsourcing measure. Our data allow us to distinguish, at the patent level, whether the patent was filed by in-house attorneys or whether the applicant was represented by an external patent professional. This information is not readily available in the EPOLINE data, but has been generated by closely examining information on each applicant's and representative's identity and address in the database. If European applicants choose to represent themselves before the EPO the 'representative' field in the database remains empty and we code the cases as internally processed. If, however, the database contains information on the representative's identity, further checks are carried out, as these cases may not necessarily all be cases in which filing is truly outsourced; for example, in some instances an in-house IP department is listed as the applicant's representative. We identify these ambiguous cases by comparing the information contained in the 'applicant name' and 'representative name' fields of the database. If the applicant's and representative's names or addresses are (partly) identical, we assume that these filings were processed internally. Finally, we cross-check this classification by searching the representative information for strings typically associated with internal patent departments (such as 'IP department,' 'Patentabteilung') and external law firms (such as '& partner'). We finally aggregate the patent-level observations and compute the average share of outsourced patent applications on the firm-year level to be included in our regressions.

Further, we construct the *upstream similarity* variable by measuring the technological similarity of internally and externally processed patent applications by comparing how similarly they are distributed across 30 different technology areas (see Footnote 12). We define two distribution vectors $\boldsymbol{f}$, the elements of which are the shares of patents a firm files in each of the 30 different technology classes: one vector contains the patents that have been processed internally ($\boldsymbol{f_i}$), and the other those that have been filed externally ($\boldsymbol{f_o}$). A measure of proximity is then given by the coefficient of uncentered correlation, or $Upstream\ Similarity = \frac{f_i f_o'}{((f_i f_i')(f_o f_o'))^{1/2}}$. If the distribution of internal and external applications coincides, similarity will take a value of 1; if there is no overlap in technology areas, the value will be 0.[17] The uncentered coefficient of

[16] Although identifying litigation targets may sound trivial at first sight, the task of spotting potentially conflicting patent rights is not straightforward given the large number of annual patent applications, which increased to more than 140,000 patent applications at the EPO in 2000, see http://documents.epo.org/projects/babylon/eponet.nsf/0/3105be5bb9ac7dfbc12574240051c6e2/$FILE/facts_figures_01.pdf, p. 14, last accessed on 24 August 2009.

[17] Zero values for internally or externally sourced production require special attention, as the uncentered correlation coefficient is not defined for these values. In these cases, when firms sourced zero products internally or externally, we set the value for upstream similarity to 0 (which comes close to a smoothing for the corner solutions). Also, it appears noteworthy that the *upstream similarity* measure is susceptible to concentration effects of a firm's patenting activities across technological areas. By controlling for such concentration effects (see our description

DOI: 10.1002/smj

correlation has been widely used in measuring the technological proximity of patents (Jaffe, 1988). Our similarity index is a refined measure of the extent to which firms are making and buying the same things—which is the case if the similarity between internally and externally processed patents in terms of technology is high.

Control variables

In order to truly capture vertical knowledge complementarities in our research, we seek to eliminate the most obvious alternative explanations that could corroborate our findings. Also, firms may decide to outsource their patent filing services under a series of constraints that could disguise the relationships we attempt to examine. We control for alternative explanations using a variety of additional variables computed at the level of the firm-year.

One of the most obvious variables that could confound our results is the firm-year rate of *outsourcing downstream activities*. In order to exclude that our findings are spuriously driven by this factor, we compute a variable that proxies the degree to which firms outsource their litigation activities (at least in part) to external representatives.[18]

Moreover, a firm's downstream performance may be affected by other factors in addition to the forgetting of related prior upstream knowledge. Elaborating on but partly deviating from earlier work (Brusoni *et al*., 2001), we argue that the more narrowly a firm conceives of its market upstream in times of change, the less likely it is to acquire knowledge relevant for downstream activities—irrespective of the firm's degree of outsourcing. More specifically, we suggest that the more concentrated a firm's engagement in patenting across few technology areas (whether through internal or external patent filing), the more likely it will fail to learn about emerging competitors (= litigation targets) in times of technological change. Such change, however, will inevitably matter in our data, which stretches over a period of approximately 20 years. We thus compute an additional control variable to describe the firm's *upstream focus* in each given year by creating a vector that captures the firm's concentration of total annual filing activity (internal *and* external filings) across the 30 different technology areas mentioned above. *Upstream focus* is measured as $Upstream\ focus_{it} = \sum_{k=1}^{30} s_{ikt}^2$, where s_{ik} is the percentage of patent applications filed in technology area k (out of 30 different technology areas) for a given firm i in year t. $Upstream\ focus_{it}$, based on 30 technology areas, is bound between 0.03 and 1. It will take on low values if a firm is active in a wide range of different technological fields. It will be high if most applications filed are concentrated in a few fields. Upstream focus serves as a proxy of how prone firms appear to miss out on learning about changes in their technological environment in general, given their overall patenting behavior.[19]

We control for industry-specific differences in patenting behavior by calculating the share of patents a firm holds in complex industries (*complex industry share*). The motives to patent differ greatly across industries. At a high level, however, it is widely held that patent filing strategies vary mostly between discrete and complex industries (Cohen, Nelson, and Walsh, 2000; Kusonaki, Nonaka, and Nagata, 1998). Discrete technologies are characterized by a relatively strong product-patent link, for example in pharmaceuticals or chemistry; whereas in complex industries, products are likely to build on technologies protected by a large number of patents held by various parties. Just as firms' patent filing and litigation strategies differ between these types of industry, their outsourcing patterns might be driven by industry context. In this paper, we control for industry-specific effects by computing a firm's yearly share

of *upstream focus* below), however, we empirically eliminate undesired distortions in the measurement of *upstream similarity*.

[18] Computing this variable appears to be the best way to control for the alternative explanation that a firm's downstream performance is driven not by the degree of outsourcing upstream but by outsourcing downstream activities themselves. This being said, the nature of the variable we could compute from the given data remains afflicted with some uncertainty. For a variety of reasons, which we do not explain in detail here, the imperfection of this variable should, if at all, introduce a conservative bias to our estimations. More information is available from the authors.

[19] More sophisticated proxies could theoretically be constructed that would better link 'missed learning' to both the actual level of a firm's external sourcing and actual changes in the firm's technological environment. Constructing and focusing on such measures, however, appears to be worth a research project of its own and is clearly beyond the scope of this paper in which we merely seek to add an additional control that we do not interpret quantitatively. Finally, it is noteworthy that *upstream focus*, by capturing possible performance-related effects stemming from the concentration of a firm's patenting across technology classes, acts as a control that facilitates an unambiguous interpretation of results pertaining to *upstream similarity*.

of patent applications related to complex technologies.[20] This approach might appear simplistic, but is one of the very few ways to disentangle industry effects from firm-specific effects for those corporations that are active in various industries at the same time.[21]

We also include a variable for firms' capacity limits to hire internal staff, limits that are due to *fluctuations in the demand for upstream activities.* The need to accommodate and smooth out fluctuations in workload is a prime reason for outsourcing human-capital-intensive tasks to external contractors (Abraham and Taylor, 1996; Houseman, 2001). Moreover, unexpected (large) increases in demand for services often lead to situations in which internal departments cannot cope with workload peaks and must therefore employ external contractors. We control for such fluctuations by measuring the steadiness of a firm's demand for patenting services over time and construct a basic volatility measure of a firm's application stream, which is computed as

$$Fluctuation_{it} = \frac{\sqrt{\frac{1}{5}\sum_{j=t-4}^{j=t}\left(PA_{ij} - \sum_{j=t-4}^{j=t} PA_{ij}/5\right)^2}}{\sum_{j=t-4}^{j=t} PA_{ij}/5}.$$

With PA_{it} representing the number of patent applications filed by applicant i in year t, this is simply the standard deviation of the applications of the preceding five years normalized by the average number of applications of the preceding five years. We normalize the measure to account for the increased relevance that absolute fluctuations have for firms issuing few applications (compared with those that issue many).

We include the variable *cumulative upstream activities,* which captures the fact that firms with larger patent portfolios can, at least in our setting, use their scale in upstream activities to partly substitute for downstream engagement via a knowledge-unrelated mechanism. Here, we measure the firm's cumulative number of patent applications over time. As is well known from the literature, a firm's ability to force its competitors into cross-licensing agreements may serve as a partial substitute for proactive patent litigation (Hall and Ziedonis, 2001). The stock of total patent applications serves as a control for the firm's ability to substitute litigation. Moreover, this variable may convey a firm's cumulative learning in the field of patenting over time, which may influence the firm's organizational forgetting rate in return.

The firm's ability to generate knowledge from the process of preparing for a patent application is a function of the technological quality of the invention, partly reflected in the patent characteristics. The literature identifies a variety of bibliographic indicators that measure a patent's technological quality (Reitzig, 2004): the number of claims, the number of different states in which patent protection is being sought, whether the application was initially filed under the Patent Co-operation Treaty, whether the applicant requested accelerated examination, the number of different technology classifications a patent has been assigned by the examiner, the total number of references, as well as the share of type A and the share of type X references to previous patents a patent application contains, and the number of references to nonpatent documents. We first compute all of the aforementioned nine indicators at the patent level. We then aggregate each indicator individually to the firm-year level. The resulting nine firm-year level variables are included in our regressions to control for *content quality*.

In order to capture unobserved effects associated with institutional changes of our empirical setting over time, we include year dummies (*year fixed effects*) in all of our regressions. Thus we can control for the fact that the European patent system has experienced a large increase in patent applications over the last few decades, which has led to organizational problems at the EPO (Harhoff and Wagner, 2009) and related institutional capacity limits.

[20] As a major robustness check to this way of modeling industry-specific differences, we tested parts of our results against findings obtained using an alternative approach originally proposed by Reitzig and Puranam (2009). In essence, we went back to our raw data and normalized every patent-level observation by its yearly industry mean. We then reconstructed our dataset using the normalized raw data and replicated our regressions on the data using ordinary least squares (OLS) fixed-effect estimations. The findings we obtained were highly robust across the two different estimation techniques.

[21] The distinction of which patents are related to complex technologies is based on the technology class, provided by the patent office and listed on a patent application. Based on those categories, we classify each patent application as related to a complex or to a discrete technology by employing the classification scheme of Cohen *et al*. (2000).

Moreover, some changes in the regulatory framework have occurred that might influence firms' patenting behavior. Also, we control for country-specific differences by including dummy variables for applicants from the three largest European applicant nations at the EPO (Germany, France, and the United Kingdom) and use applicants from the remaining states as the reference group.

Model specifications and econometric issues

Optimally, we would like to relate *downstream performance* (firm-year number of oppositions/firm-year number of patent applications) as the dependent variable to the independent variables, notably *outsourcing rate of upstream activities*. However, this dependent variable would be bound between 0 and 1. This condition leaves us with a nontrivial estimation problem. The two-sided censoring of our data represents a major infringement of standard Gauss-Markov assumptions, suggesting the use of a Tobit model instead of OLS. However, to the best of our knowledge, there exists no Tobit model that could account for firms' fixed effects in a panel. Therefore, capturing unobserved heterogeneity appears crucial given the exploratory nature of this study.

Given the difficulties, we therefore approach the estimation problem differently. We employ a panel count data model based on Poisson regression, rearranging the estimation problem so that the *denominator of the downstream performance* ratio (firm-year number of patent applications) appears as an independent variable (see below for details), and the *numerator of downstream performance*, a count measure (firm-year number of oppositions), becomes our dependent variable. In this way the Poisson model[22] indirectly caters to our estimation challenge without loss of information. At the same time, however, we can now control for unobserved heterogeneity by eliminating (conditional) fixed effects (CFE). Finally, in order to exclude the CFE results suffering from a systematic selection bias,[23] we compared the results of the CFE Poisson model to those of a simple panel Poisson and checked on the robustness of our findings.[24] In the following, we describe in more detail how the panel Poisson model can cater to our estimation problem.

Following Hausman, Hall, and Griliches (1984), we choose a basic Poisson count estimation. In the Poisson model, the count variable dep_{it} for firm i in year t is assumed to follow a distribution as follows:

$$dep_{it}|\lambda_{it} \sim Poisson(\lambda_{it}) \quad (1)$$

We consider specifications of the form $\lambda_{it} = E(dep_{it}|X_{it}) = \exp(X_{it}\beta)$, where X_{it} is a vector containing our independent and control variables. Additionally, we allow λ_{it} to be a function not only of observable variables X_{it} but also of unobserved firm-specific effects. These firm-specific effects are assumed to be time invariant and might be interpreted as differences in the possession of capabilities or other factors that influence a firm's success in generating and litigating IP rights. In the following, we denote these effects as μ_i and introduce them in a multiplicative way (Cameron and Trivedi, 1998; Hausman *et al*., 1984). Consequently, our final regression model can be written as

$$dep_{it}|\lambda_{it}, \mu_i \sim Poisson(\lambda_{it}) \quad (2)$$

with

$$\lambda_{it} = E(dep_{it}|X_{it}, \mu_i) = \exp(X_{it}\beta + \mu_i). \quad (3)$$

A reformulation of Equation 3 yields the more familiar log-linear form with

$$\log(\lambda_{it}) = \log(E(dep_{it}|X_{it}, \mu_i)) = X_{it}\beta + \mu_i. \quad (4)$$

Since, initially, we are not interested in the absolute number of yearly oppositions filed by the firm but rather in the share of yearly oppositions filed relative to the total number of applications filed per year, we include the logarithm of the total number of patent applications PA_{it} to the set of independent variables. The average share of

[22] While the Poisson model is restrictive because it implies the assumption that the variance of λ_{it} is equal to its mean, we prefer this model to the more flexible negative binomial model. It is known that the negative binomial model is highly sensitive against violations of the underlying assumptions, while the Poisson model is consistent as long as the mean specification holds (Gourieroux, Monfort, and Trognon, 1984).

[23] As is well known, panel Poisson CFE models drop observations for those firms whose independent variables do not change over time. This could introduce a systematic selection bias. By comparing the panel Poisson CFE results to those of a simple panel Poisson model, we obtain an understanding of the magnitude of this effect.

[24] More information is available from the authors upon request.

oppositions relative to patent applications $\widetilde{\lambda}_{it}$ can then be expressed as

$$\begin{aligned}\widetilde{\lambda}_{it} \cdot PA_{it}^{\gamma} &= \widetilde{\lambda}_{it} \cdot \exp(\gamma \log PA_{it}) \\ &= \exp(X_{it}\beta + \mu_i + \gamma \log PA_{it}). \qquad (5)\end{aligned}$$

Note that $\widetilde{\lambda}_{it} \cdot PA_{it}$ in Equation 5 can be interpreted as the count of oppositions filed by firm i in year t. γ in Equation 5 is a measure of the returns to scale in patent filing. Whereas a unit change in a variable x_k leads to a change in the conditional mean by the amount $E(dep_{it}|X_{it}, \mu_{it}) \times \beta_k$ and, therefore, to a proportionate change in $E(dep_{it}|X_{it}, \mu_{it})$ by β_k, the coefficient γ has to be interpreted as the elasticity of the number of firm-year oppositions filed with regard to PA_{it}. If the estimated value of γ is significantly different from 1, the share of filed oppositions is not proportional to the yearly number of patent applications. In particular, a coefficient larger than 1 could be interpreted as evidence of positive returns to scale from patent filing activities. The estimated coefficient β can be interpreted as the effect of the independent and control variables on the share of filed oppositions relative to patent applications because the total number of patent applications is included in the regressions.

EMPIRICAL RESULTS

Table 1 contains summary statistics for the major variables used in the following analyses. Table 2 provides the Pearson correlations between the major variables of interest. We find low to moderate correlations across all variables.

Our main results from the multivariate Poisson regressions are presented in Table 3. Table 3 summarizes the tests we conducted in connection with Hypotheses 1 and 2.

In Hypothesis 1, we postulated that a firm's downstream performance (as measured by the patent opposition/patent application ratio) decreases (increases) the more (less) the firm sources upstream services (patent filing services) externally. Table 3, Column A, provides empirical support for this hypothesis. The coefficient for the outsourcing rate has a negative effect on the downstream performance numerator, whereas the effect of the log (*downstream performance denominator*) is positive. The tests are carried out on our entire sample. It also indicates (coefficient of *log*(*downstream performance denominator*) <1) that firms that patent frequently are less likely than other firms to engage in downstream activity relative to the size of their patent portfolio.

Also, Column B of Table 3 confirms Hypothesis 2. Although the degree to which a firm outsources upstream (patent filing) services reduces its downstream (litigation) performance, similarity between internal and external upstream activities can, as a main effect, offset this total knowledge loss to some extent, as expected.

DISCUSSION AND CONCLUSIONS

Little research exists on how a firm's performance is influenced by how it draws its boundaries

Table 1. Descriptive statistics

Variable	Obs	Mean	Std. dev.	Min	Max
Downstream performance numerator[a]	7960	1.75	5.79	0	113
Downstream performance denominator[a]	7960	23.78	50.89	1	661
Outsourcing rate of upstream activities[a]	7960	0.48	0.45	0	1
Upstream similarity[a]	7960	0.20	0.34	0	1
Cumulative upstream activities[a]	7960	238.11	654.32	1	9,757
Upstream focus	7960	0.47	0.28	0.06	1
Fluctuations in the demand for upstream activities[a]	7960	6.81	12.98	0	363
Complex industry share[a]	7960	0.70	0.46	0	1
Outsourcing rate of downstream activities[a]	7960	0.68	0.19	0	1
German applicant (0/1)	7960	0.42		0	1
UK applicant (0/1)	7960	0.12		0	1
French applicant (0/1)	7960	0.19		0	1

[a] Variable computed at the firm-year level.

DOI: 10.1002/smj

Table 2. Correlation of major variables

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1)	Downstream performance numerator	1.00								
(2)	Downstream performance denominator (log)	0.40	1.00							
(3)	Outsourcing rate of upstream activities	−0.11	−0.18	1.00						
(4)	Upstream similarity	0.16	0.36	−0.14	1.00					
(5)	Cumulative upstream activities	0.60	0.43	−0.11	0.21	1.00				
(6)	Upstream focus	−0.18	−0.50	0.08	−0.13	−0.20	1.00			
(7)	Fluctuations in the demand for upstream activities	0.09	0.14	−0.05	0.13	0.30	−0.04	1.00		
(8)	Complex industry share	0.08	0.24	−0.01	0.05	0.09	−0.19	0.04	1.00	
(9)	Outsourcing rate of downstream activities	−0.09	−0.07	0.24	0.00	−0.08	0.00	−0.05	−0.05	1.00

(Leiblein, Reuer, and Dalsace, 2002; Rothaermel *et al.*, 2006). Even less research focuses explicitly on the link between firm performance and knowledge complementarities of vertically related activities (Brusoni *et al.*, 2001; Jacobides and Billinger, 2006). Large-scale evidence for these recent theoretical advances is lacking entirely. Finally, the possibility of firms losing existing knowledge through forgetting by outsourcing vertically related activities—here termed 'hidden outsourcing costs'—has been neither explicated theoretically nor tested empirically. Our results provide an initial attempt to close these knowledge gaps and to reinforce and refine the recent theoretical understanding about the link between firm boundaries and firm performance.

The first aim of the paper was to put recent theoretical conjectures on vertical knowledge complementarities, gained from case-based work, on more solid empirical ground. Despite the usual limitations of large-scale tests as well as limitations specific to our design,[25] our findings as illustrated in Column A of Table 3 quite clearly indicate that firms' downstream performance (in our setting, their ability to detect competitors for proactive litigation purposes) decreases (increases) the more (less) they outsource the preceding vertically related upstream services (in our setting, patent filings). A shift from complete integration to complete outsourcing upstream leads to a 27 percent drop in litigation performance in our data. We interpret these findings to confirm the relationships we hypothesized in Hypothesis 1. We can exclude that our results are a spurious effect driven by the firm's outsourcing of downstream (litigation) activities themselves. Moreover, our results are robust across various specifications that include fewer controls. The systematic selection bias introduced by the CFE Poisson model is desirable, given that we want to look at firms for which knowledge complementarity matters.[26] Naturally, our findings, although expressed in general form, are ultimately obtained within a particular context. Hence, we seek to control for as many context-specific variables as possible to ensure that our central claim appears to be independent of the

[25] We mention two potential caveats in particular. First, our measure of institutional capacity constraints at the EPO is imperfect in that we capture only aggregate capacity constraints across industries. As is known from the literature (Harhoff and Wagner, 2009), however, these capacity constraints may play out in distinctly different ways across industries. Thus, our measure may imperfectly capture the true capacity constraints for highly technologically specialized firms. Second, despite our efforts to exclude endogeneity distortions in our design, we cannot rule out that more complex selection effects are at work in our sample that we cannot capture with the present estimations. We deem the problems minor given the various cross-checks we ran, however, particularly given the confirmatory results of alternative dummy variable estimations (Parmigiani, 2007). More information is available from the authors upon request.

[26] The conditional fixed-effect model specifically drops those firm-year observations when downstream success rates (opposition) are constant over time. In our sample, the vast majority of these observations are cases in which firms file absolutely no oppositions. These firms are to be considered structurally different from the rest of the sample in that they appear to attribute no value to downstream activities at all; hence, knowledge complementarities between upstream and downstream activities should consequently not matter to them either, and it is desirable to exclude them from the estimations. This being said, results are quite robust across specifications even if we do not account for conditional fixed effects (results available upon request).

Table 3. **Main results from the multivariate Poisson regressions**

	Downstream performance numerator	
	A	B
Downstream performance denominator (log)	0.361*** [0.018]	0.351*** [0.018]
Outsourcing rate of upstream activities	−0.312*** [0.059]	−0.318*** [0.059]
Upstream similarity	—	0.074** [0.037]
Upstream focus	−0.675*** [0.090]	−0.689*** [0.090]
Share of outsourced downstream activities	0.583*** [0.107]	0.588*** [0.107]
Fluctuations in the demand for upstream activities	−0.003*** [0.001]	−0.003*** [0.001]
Cumulative upstream activities (coefficient × 1000)	−0.139*** [0.012]	−0.142*** [0.011]
Complex industry share	−0.051 [0.033]	−0.051 [0.033]
German applicant	0.285*** [0.058]	0.285*** [0.058]
UK applicant	−0.058 [0.085]	−0.058 [0.085]
French applicant	0.072 [0.100]	0.072 [0.100]
Year dummies	YES	YES
Content quality	YES	YES
Observations	6154	6154
Number of firms	373	373
Log-likelihood	−7000.5	−6998.6

Downstream performance as a function of upstream outsourcing rate (conditional fixed effects Poisson models, standard errors in parentheses). Dependent variable split into numerator (= number of patent oppositions) and denominator (= number of patent applications). Denominator enters as explanatory variable in rearranged Poisson model.
Significantly different from 0 at the * 10% level, ** 5% level, *** 1% level based on two-sided t-tests.

peculiarities of our empirical testing ground. These control variables show the expected signs for the most part; in particular, as we expected, the more concentrated a firm's overall patenting engagement across few technology categories, the more it misses out on learning about potential competitors, which leads, in turn, to lower litigation activity downstream.[27] Moreover, we find confirmatory evidence that firms with less downstream dependence (i.e., firms with larger patent portfolios as reflected by their cumulative upstream activities), an important construct in empirical settings related to patent litigation, need to rely less on downstream activities (and hence file fewer oppositions). It is likely that the variable does not, as one could have thought *ex ante*, reflect prior learning on the firm's part because it does not offset the knowledge losses from outsourcing.[28] National idiosyncrasies in using specifically European litigation procedures exhibit an effect that is independent of firm-specific attributes. Also, firms that apply for more patent applications in a given year suffer stronger performance dips than do firms that apply for fewer patents. Finally, outsourcing downstream activities to litigation experts increases downstream performance in Model 3A (Table 3, Column A)—which is plausible but not of major interest in this study as it does not substitute the effect of upstream outsourcing.

[27] A knowledge-unrelated mechanism for this finding would involve the assumption that focused firms have less need for downstream litigation. We cannot rule out that this mechanism plays a role, but it does not strike us as particularly relevant.

[28] At least, no corresponding main effect is visible.

The second aim of the paper was to refine our understanding of the differences between vertical integration benefits and vertical outsourcing costs, if they exist. Particularly, we wondered whether firms simply *forget* existing knowledge through outsourcing. Our empirical setting appears highly suited to study such hidden outsourcing costs. If they exist at all, they should prevail in situations when the outsourcing rate increases, just as is the case in our data (see Figure 2).[29] Moreover, our setting appears suited for our tests for a second reason. Namely, hidden outsourcing costs, if at work, should be the only outsourcing-related costs within the design we have chosen, as other outsourcing-related costs appear to be negligible in our setting.[30] In turn, however, this means that our data lends itself particularly well to the large-scale examination of hidden outsourcing costs, as these costs may explain to a large extent why firms would still not outsource all of their upstream activities.

As for Hypothesis 1, the context-specificity of our data and the reference to underlying and partially untested assumptions[31] require us to interpret our findings with respect to Hypothesis 2 cautiously. Within these confines, however, we find empirical support for our second hypothesis, although it is somewhat less stable than for Hypothesis 1.[32] To summarize, our results suggest that hidden outsourcing costs do exist and that they differ conceptually from their seeming flipside of vertical integration benefit. In more detail, we find that sourcing similar upstream services (patent applications) in-house and from external suppliers increases downstream performance, all else being equal. As theorized, our interpretation is that similar internal and external sourcing patterns upstream lead to an increased retention of knowledge inside the corporation that will compensate for some of the vertical outsourcing costs through forgetting.[33] Notably, this also implies that forgetting through outsourcing is one of the mechanisms at work in our data.

That said, the forgetting of prior knowledge and the missing of learning opportunities, a distinctly different mechanism driving vertical knowledge costs, may jointly account for the overall vertical knowledge losses through outsourcing that we detect. Although we cannot provide conclusive evidence that firms did miss learning opportunities through excessive *external sourcing*, it does appear as though, *more generally* speaking, failing to acquire downstream-related knowledge from upstream activities harmed firms in our sample. As in Model 3A, more narrowly focused firms spot fewer competitors than other firms also in Model 3B.

Despite its limits, we believe that the analysis in this paper contributes to the literature on innovation management as well as organization design more broadly. Recently, the literature on firms' strategic use of IP has begun to address questions of organizational structure (Reitzig and Puranam, 2009; Somaya *et al*., 2007; Wagner, 2007). This paper extends the current stream of research by refining the picture of specialization advantages and coordination needs (Reitzig and Puranam, 2009; Somaya *et al*., 2007) in the IP value chain

[29] The overall average trend of more patent filing services being sourced from external sources over time does not apply to all firms in our sample. In order to exclude that our results for *upstream similarity* would misleadingly be driven by the *learning* of the subsample of firms that upscale their vertical activities over time, we ran a series of robustness checks. More information is available from the authors upon request.

[30] Notably, within the IP value chain the outsourcing of upstream (patent filing) services is associated with substantial efficiency gains for firms as long as performance is measured solely along a dimension of upstream performance (patent grant success) (see Reitzig and Wagner, 2009).

[31] We are unable to test whether all of our core assumptions are being fulfilled. We therefore chose a different approach to assess the distortions arising from a potential infringement of one or several assumptions we make. Namely, we wondered whether any potential combination of suppositional infringements could cause our similarity to pick up knowledge-related effects other than forgetting through outsourcing. Barring any such visible combination of infringements, we conclude that any violation of our assumptions will, at best, introduce a conservative bias into our estimations. Additional note: characterizing forgetting-related outsourcing costs by alternative measures that draw on different assumptions led to counterintuitive results. More results are available from the authors.

[32] We ran a series of robustness checks in an attempt to exclude that the findings for Hypothesis 2 are driven by specificities of our data. Not all robustness checks support Hypothesis 2 equally well; however, we can confidently exclude that our results are regression artifacts that are spuriously driven by limitations of our data, notably truncation. More information is available from the authors upon request.

[33] Looking at the marginal effect (percentage change in the number of expected oppositions filed), retaining similar knowledge can mitigate only about seven percent of the knowledge loss in litigation caused through the outsourcing of patent filing, which we calculate to be 27 percent in Model 3B when going from complete upstream integration to complete upstream outsourcing.

in that it reflects on the boundaries of the innovative firm (Cassiman and Veugelers, 2006; Pisano, 1990; Wagner, 2007). Our research suggests that knowledge complementarities exist not only within the domain of IP generation (e.g., R&D see Cassiman and Veugelers, 2006) but also across the spectrum of vertically related IP-protection tasks. Firms appear to face a trade-off between leveraging the direct benefits from outsourcing patent filing activities to external markets and losing litigation-related knowledge. How the elucidation of this trade-off depends on environmental constraints, particularly the strength of the appropriability regime (Cohen *et al*., 2000), remains an open question to be addressed in future studies.

Our findings are also likely to be relevant to a broader audience of organizational scholars. At a more general level we consider our findings to be an empirical confirmation of recent theoretical advances in the knowledge-based theory of the firm. Building on arguments of knowledge complementarity across (vertically related) activities, and adding the perspective of organizational memory (Cyert and March, 1963; Huber, 1991; Levitt and March, 1988; March and Simon, 1958) to the discussion, we propose and empirically test the existence of hidden outsourcing costs. Our current results confirm the existence of knowledge-based complementarities between adjacent activities in firms' value chains and represent the first large-scale test for these conjectures. Our results also indicate that understanding the outsourcing-driven knowledge-loss phenomenon along the value chain requires applying lenses that are more sophisticated than simple cost-benefit analyses, which fail to consider the particularities of organizational learning versus forgetting. Whereas standard economics would merely equate dynamic vertical integration benefits with the opportunity costs of knowledge losses from the outsourcing of vertically related activities, we propose and find differences between the processes of forgetting and acquiring knowledge. Most importantly, a cautious normative interpretation of our findings suggests that vertically integrated firms should avoid taking the (short-term) bait of outsourcing without considering its potential impact on their performance in preceding or subsequent steps of the value chain. Such considerations may lead firms to outsource their core activities in a seemingly inefficient way in that they try to source similar activities both internally and externally to avoid knowledge losses.

Within the confines of our specific empirical design, we can make some additional interesting and novel quantitative observations. Although retaining knowledge through targeted upstream outsourcing (= high similarity between internally and externally sourced upstream production) reduces vertical outsourcing costs, it cannot offset them entirely. Furthermore, vertical outsourcing costs are subject to the volume of total upstream activities sourced by the firm, stressing the importance of implementing routines and procedures, particularly in large firms that can counterbalance the undesired evaporation of knowledge that accompanies outsourcing. Finally, it appears from specifications not presented in this paper as though hidden outsourcing did not rise nonlinearly in the level of upstream outsourcing. One future avenue that appears to be particularly promising to us is the study of further differences between integration benefits through learning and forgetting-related outsourcing costs. The sparse literature on organizational learning versus forgetting suggests that, more often than not, forgetting through outsourcing may happen more rapidly than learning through integration. This is a consequence of the asymmetries that occur when filling and emptying the 'knowledge retention bins' (Olivera, 2000; Walsh and Ungson, 1991). Little is known about firms' memory and forgetting as opposed to their learning (Shafer, Nembhard, and Uzumeri, 2001). The existing experimental and empirical evidence suggests several reasons, however, why firms should forget faster than they learn, the most universally applicable being personnel turnover (Benkhard, 2000; Darr *et al*., 1995;). Such turnover should, all else being equal, always inhibit a firm's learning and accelerate its forgetting. Data in addition to ours are needed, however, to empirically investigate these questions.

ACKNOWLEDGEMENTS

The authors thank Associate Editor Rudi K.F. Bresser as well as two anonymous referees for their guidance and support during the submission and revision process. Markus Reitzig also thanks his colleagues Freek Vermeulen, Michael Jacobides, Brandon Lee, and particularly Phanish Puranam for extensive discussions on this paper. The authors are responsible for all remaining errors.

REFERENCES

Abraham KG, Taylor SK. 1996. Firms' use of outside contractors: theory and evidence. *Journal of Labor Economics* **16**(3): 395–424.

Argote L. 1996. Organizational learning curves: persistence, transfer, turnover. *International Journal of Technology Management*, Special Issue on Unlearning and Learning for Technological Innovation **11**: 759–769.

Argyres NS, Liebeskind JP. 1999. Contractual commitments, bargaining power, and governance inseparability: incorporating history into transaction cost theory. *Academy of Management Review* **24**(1): 49–63.

Benkard C. 2000. Learning and forgetting: the dynamics of aircraft production. *American Economic Review* **90**(4): 1034–1054.

Brusoni S, Prencipe A, Pavitt K. 2001. Knowledge specialization, organization coupling, and the boundaries of the firm: why do firms know more than they make? *Administrative Science Quarterly* **46**(4): 597–621.

Cameron AC, Trivedi PK. 1998. *Regression Analysis of Count Data*. Cambridge University Press: Cambridge, MA.

Campbell-Kelly M. 1996. Information technology and organizational change in the British census, 1801–1911. *Information Systems Research* **7**(1): 22–36.

Cassiman B, Veugelers R. 2006. In search of complementarity in innovation strategy: internal R&D and external knowledge acquisition. *Management Science* **52**(1): 68–82.

Coase R. 1937. The nature of the firm. *Economica* **4**(16): 386–405.

Cohen WM, Levinthal DA. 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* **25**: 128–152.

Cohen WM, Nelson RR, Walsh JP. 2000. Protecting their intellectual assets: appropriability conditions and why U.S. manufacturing firms patent (or not). NBER Working Paper WP 7552, National Bureau of Economic Research, Cambridge, MA.

Cyert RM, March JG. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall: Englewood Cliffs, NJ.

Darr ED, Argote A, Epple. D. 1995. The acquisition, transfer, and depreciation of knowledge in service organizations: productivity in franchises. *Management Science* **41**(11): 1750–1762.

Dyer JH, Hatch NW. 2006. Relation-specific capabilities and barriers to knowledge transfers: creating advantage through network relationships. *Strategic Management Journal* **27**(8): 701–719.

Dyer JH, Singh H. 1998. The relational view: cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review* **23**(4): 660–679.

EPO. 2000. *European Patent Convention.* Part VII Common provisions, Chapter III Representation. European Patent Office, Munich, Germany.

Gourieroux C, Monfort A, Trognon A. 1984. Pseudo-maximum likelihood methods: applications to Poisson models. *Econometrica* **52**(3): 701–720.

Grant RM. 1996. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, Winter Special Issue **17**: 109–122.

Gulati R, Lawrence PR, Puranam P. 2005. Adaptation in vertical relationships: beyond incentive conflict. *Strategic Management Journal* **26**(5): 415–440.

Hall BH, Ziedonis RH. 2001. The patent paradox revisited: an empirical study of patenting in the U.S. semiconductor industry 1979–1995. *RAND Journal of Economics* **32**(1): 101–128.

Hargadon A, Sutton RI. 1997. Technology brokering and innovation in a product development firm. *Administrative Science Quarterly* **42**: 716–749.

Harhoff D, Reitzig M. 2004. Determinants of opposition against EPO patent grants: the case of pharmaceuticals and biotechnology. *International Journal of Industrial Organization* **22**(4): 443–480.

Harhoff D, Wagner S. 2009. The duration of patent examination at the European patent office. *Management Science* **55**(12): 1969–1984.

Hatch NW, Dyer JH. 2004. Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal* **25**(12): 1155–1178.

Hausman J, Hall BH, Griliches Z. 1984. Econometric models for count data with an application to the patents-R&D relationship. *Econometrica* **52**(4): 909–938.

He D, Nickerson JA. 2006. Why do firms make and buy? Efficiency, appropriability and competition in the trucking industry. *Strategic Organization* **4**(1): 43–69.

Helper S, MacDuffie JP, Sabel C. 2000. Pragmatic collaborations: advancing knowledge while controlling opportunism. *Industrial & Corporate Change* **9**(3): 443–488.

Hinze S, Reiss T, Schmoch U. 1997. Statistical analysis on the distance between fields of technology. Innovation Systems and European Integration (ISE): Sub-Project 3.2.5: Technological entry: diversification vs. new innovators. Report submitted to the European Commission.

Houseman SN. 2001. Why employers use flexible staffing arrangements: evidence from an establishment survey. *Industrial and Labor Relations Review* **55**(1): 149–170.

Huber GP. 1991. Organizational learning: the contributing process and the literatures. *Organization Science* **2**(1): 88–115.

Jacobides MG, Billinger S. 2006. Designing the boundaries of the firm: from 'make, buy, or ally' to the dynamic benefits of vertical architecture. *Organization Science* **17**(2): 249–261.

Jaffe AB. 1988. Demand and supply influences in R&D intensity and productivity growth. *Review of Economics and Statistics* **70**(3): 431–437.

Knight HJ. 1996. *Patent Strategy for Researchers and Research Managers*. Wiley: Chichester, UK.

Kogut B, Zander U. 1992. What do firms do? Coordination, identity, and learning. *Organization Science* **7**(5): 502–518.

Kusonaki K, Nonaka I, Nagata A. 1998. Organizational capabilities in product development of Japanese firms. *Organization Science* **9**(6): 699–718.

Leiblein MJ, Reuer JJ, Dalsace F. 2002. Do make or buy decisions matter? The influence of organizational governance on technological performance. *Strategic Management Journal* **23**(9): 817–833.

Levitt B, March JG. 1988. Organizational learning. *Annual Review of Sociology* **14**: 319–340.

March JG, Simon HA. 1958. *Organizations*. Wiley: New York.

Nelson RR, Winter SG. 1982. *An Evolutionary Theory of Economic Change*. Belknap Press: Boston, MA.

O'Connell D. 2008. *Inside the Patent Factory: The Essential Reference for Effective and Efficient Management of Patent Creation*. Wiley: Chichester, UK.

OECD. 1994. *The measurement of science and technology indicators—using patent data as science and technology indicators*. Technical report. Organisation for Economic Co-operation and Development: Paris, France.

Olivera F. 2000. Memory systems in organizations: an empirical investigation of mechanisms for knowledge collection, storage and access. *Journal of Management Studies* **37**(6): 811–832.

Olivera F, Argote L. 1999. Organizational learning and CORE processes in new product development. In *Shared Cognition in Organizations: The Management of Knowledge*, Thompson IL, Levine JM, Messick DM (eds). Laurence Erlbaum Associates: Mahwah, NJ; 297–326.

Parise S, Cross R, Davenport TH. 2006. Strategies for preventing a knowledge-loss crisis. *Sloan Management Review*. **47**(4): 31–38.

Parmigiani AE. 2007. Why do firms both make and buy? An investigation of concurrent sourcing. *Strategic Management Journal* **28**(3): 285–311.

Pisano GP. 1990. The R&D boundaries of the firm: an empirical analysis. *Administrative Science Quarterly* **35**: 153–176.

Pisano G. 1994. Knowledge, integration, and the locus of learning: an empirical analysis of process development. *Strategic Management Journal*, Winter Special Issue **15**: 85–100.

Puranam P, Gulati R, Bhattacharya S. 2008. How much to make and how much to buy? Explaining plural sourcing strategies. Working paper, London Business School, London, UK.

Reitzig M. 2004. Improving patent valuation methods for management—validating new indicators by analyzing application rationales. *Research Policy* **33**(6–7): 939–957.

Reitzig M. 2007. How executives can enhance IP strategy and performance. *Sloan Management Review* **48**(3): 37–43.

Reitzig M, Puranam P. 2009. Value appropriation as an organizational capability: the case of IP protection through patents. *Strategic Management Journal* **30**(7): 765–789.

Reitzig M, Wagner S. 2009. Patent grant success as a function of patent filing outsourcing. Mimeo, London Business School, London, UK.

Rothaermel FT, Hitt MA, Jobe LA. 2006. Balancing vertical integration and strategic outsourcing: effects on product portfolio, product success, and firm performance. *Strategic Management Journal* **27**(11): 1033–1056.

Santos FM, Eisenhardt KM. 2005. Organizational boundaries and theories of organization. *Organization Science* **16**(5): 491–508.

Schilling MA, Vidal P, Polyhardt RE, Marangoni A. 2003. Learning by doing *something else*: variation, relatedness, and the learning curve. *Management Science* **49**(1): 39–56.

Shafer SM, Nembhard DA, Uzumeri MV. 2001. The effects of worker learning, forgetting, and heterogeneity on assembly productivity. *Management Science* **47**(12): 1639–1653.

Simon HA. 1991. Bounded rationality and organizational learning. *Organization Science* **2**(1): 125–134.

Somaya D, Williamson IO, Zhang X. 2007. Combining patent law expertise with R&D for patenting performance. *Organization Science* **18**(6): 922–937.

Sorenson O, Sorensen JB. 2001. Finding the right mix: franchising, organizational learning, and chain performance. *Strategic Management Journal*, June–July Special Issue **22**: 713–724.

Stinchcombe AL. 1990. *Information and Organizations*. University of California Press: Berkeley, CA.

Wagner S. 2007. Make or buy decisions in patent-related services. *Zeitschrift für Betriebswirtschaft (ZfB)* **4**: 47–68.

Walsh JP, Ungson GR. 1991. Organizational memory. *Academy of Management Review* **16**: 57–90.

Williamson OE. 1971. The vertical integration of production: market failure considerations. *American Economic Review* **61**(2): 112–123.

Williamson OE. 1985. *The Economic Institutions of Capitalism*. Free Press: New York.