American Economic Review 2014, 104(7): 1979–2013
http://dx.doi.org/10.1257/aer.104.7.1979

Patent Laws, Product Life-Cycle Lengths, and Multinational Activity†

By L. Kamran Bilir*

Do intellectual property rights influence multinationals' manufacturing location decisions? My theoretical model indicates that countries with strong patent laws attract multinational activity, but only in sectors with relatively long product life cycles. By contrast, firms with short life-cycle technologies are insensitive, because offshore imitation is less likely to succeed before obsolescence. I document strong empirical regularities consistent with the model using a panel dataset on the global operations of US-based multinational firms and a new measure of product obsolescence. Moreover, my identification strategy allows me to isolate the causal effect of patent laws on multinational activity. (*JEL* D92, F23, K11, L60, O34, R32)

Multinational corporations are among the most innovation-intensive firms and account for the substantial majority of international transactions.[1] To an increasing extent, these transactions involve proprietary technologies transferred within the firm from multinational parents to their foreign affiliates.[2] But in choosing where to establish foreign affiliates and deploy these proprietary technologies, multinational firms face a trade-off, because countries with attractive input costs often lack strong protection for intellectual property.

This trade-off does not, however, affect firms in all sectors equally. Consider, for example, the experience of firms in the electronics, hard-disk drive, and solar cell industries. Electronics and hard-disk drive firms tend to produce even their newest products in countries with weak intellectual property institutions.[3] By contrast,

* Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (e-mail: kbilir@ssc.wisc.edu). The statistical analysis of firm-level data on US multinational companies was conducted at the Bureau of Economic Analysis, US Department of Commerce under arrangements that maintain legal confidentiality requirements. The views expressed are those of the author and do not reflect official positions of the US Department of Commerce. The Bureau of Economic Analysis has reviewed this paper prior to release to ensure that data confidentiality is not unintentionally compromised. The author has no relevant or material financial interests that relate to the research described in this paper. I thank Kenneth Arrow, Nick Bloom, Kalina Manova, Petra Moser, and Bob Staiger for their invaluable guidance. I also thank Dietmar Harhoff, Ken Hendricks, Mark Lemley, Ray Mataloni, Eduardo Morales, Mitch Polinsky, Stanley Watt, and Bill Zeile for insightful conversations, and seminar participants at Boston College, Calgary, Chicago Fed, Columbia, Dartmouth Tuck, HBS, Kellogg, Maryland, NBER, Princeton, Stanford, UCLA, Wharton, and Wisconsin, and three anonymous referees for very helpful comments. Financial support from the John M. Olin Foundation and the Stanford Institute for Economic Policy Research Leonard W. Ely and Shirley R. Ely Fellowship is gratefully acknowledged.

† Go to http://dx.doi.org/10.1257/aer.104.7.1979 to visit the article page for additional materials and author disclosure statement(s).

[1] Doms and Jensen (1998); UNCTAD (2005); Criscuolo, Haskel, and Slaughter (2010).

[2] National Science Board (2010). See also Keller and Yeaple (2013).

[3] Hard-disk drive firms Seagate and Western Digital maintain manufacturing facilities in Thailand, for example, where patent institutions consistently fall between the tenth and twentieth percentiles across countries during 1985–2005; Seagate has pursued an aggressive offshoring strategy for decades (Igami 2010).

Solar Junction, a US firm that has developed high-efficiency solar cells, performs its proprietary manufacturing activity in the United States, pointing to concerns about offshore imitation risk as a major factor behind the decision. Solar Junction cites the durability of the intellectual property associated with its current products—which is long-lived compared with that embedded in electronics and hard-disk products—as a key underlying cause for its sensitivity to intellectual property laws.[4]

This article provides evidence that multinational firms' sensitivity to host-country intellectual property protection is determined by the length of product life cycles. I quantify the effects of patent laws and product life-cycle lengths on US firms' global operations across 37 industries and 92 countries during 1982–2004. My results indicate that intellectual property rights attract systematically higher levels of multinational activity in sectors with long product life cycles. I show that these patterns in the data are consistent with a global production model in which innovating firms in the North face imitation risk in the South. Firms with short life-cycle products are less sensitive to intellectual property rights in the model, because offshore imitation is less likely to succeed before obsolescence. I implement these findings to empirically isolate the causal effect of patent laws on multinational activity.

In the model, innovating firms make production location decisions for products with industry-specific life-cycle lengths that are technologically determined. At every moment, firms choose whether to locate manufacturing in the North or in the South. While firms initially enjoy higher profits when manufacturing in the South, patents there are imperfectly protected compared with in the North. This affects location decisions, because manufacturing requires the use of proprietary knowledge; the act of manufacturing exposes this knowledge to local entrepreneurs, enabling imitation to arise where manufacturing occurs. Importantly, the risks and expected losses associated with imitation depend on both the quality of local patent laws and on products' remaining economic lifetimes. The sourcing trade-off thus evolves over the product life cycle.

I show that firms follow a sector-invariant optimal location rule, whereby production moves to the South when products reach a critical time-to-obsolescence cutoff. Improvements to patent protection in the South increase this cutoff. Because products with lifetimes shorter than the cutoff are always manufactured in the South, patent reforms have no effect on location decisions in fast-turnover industries. By contrast, the manufacture of longer-lived products is offshored to the South earlier in the life cycle following the reform, increasing multinational activity in these industries. Moreover, the industry-level response to patent reforms is a nonmonotonic function of product life-cycle lengths and is most pronounced in intermediate life-cycle length sectors. Intuitively, this is because the increase in the offshoring cutoff affects the manufacturing location only for marginal product varieties, the measure of which is highest in sectors with intermediate product lifetimes. These comparative statics also hold in the cross-section of countries with different levels of patent protection. Finally, these effects are concentrated among relatively less-productive firms.

For my empirical analysis, I develop a new index of product life-cycle lengths using information in patent records from the NBER US Patent Citation Data File.

[4] Based on conversations with senior management at Solar Junction.

The index is based on the idea that the duration of citations received by patents reflects the lifetime of technologies embedded in products (e.g., Narin and Olivastro 1993); my approach is consistent with evidence that innovative ideas tend to span several successive versions over the product life cycle as in Klepper (1996). Within a simple theory of sequential innovation, I show that a specific industry-level index—the mean forward citation lag (or age)—captures variation in average product lifetimes. I use this as my primary measure; I also compare my results with estimates based on alternative proxies, including the Broda and Weinstein (2010) product turnover index. Host-country patent laws are measured with a widely used index of patent protection (Ginarte and Park 1997; Park 2008), and I measure US firms' global location decisions using data from the Bureau of Economic Analysis.

My empirical analysis indicates that, relative to short–life cycle sectors, firms in sectors with long product life cycles offshore a significantly smaller share of manufacturing activity; this is consistent with firms choosing to begin offshoring longer-lived products systematically later in the product life cycle. I also find that levels of affiliate sales, assets, and employment by industry respond to patent strength according to the industry's product life-cycle length. In particular, I find that sensitivity to patent protection is nonmonotonic in product life-cycle lengths: it is low in short life-cycle sectors and high in long life-cycle sectors, with the largest effects in sectors near the seventy-fifth percentile of the distribution. This differential sensitivity to patent laws is also economically significant: a one–standard deviation increase in measured patent protection attracts between 10 and 20 percentage points more multinational activity in the seventy-fifth percentile sector than in the tenth-percentile sector by product life-cycle lengths. Similarly, the effect at the seventy-fifth percentile exceeds that at the ninety-fifth percentile sector by between 5 and 10 percentage points.

I find that these industry-level effects reflect distinct modes of firm response: stronger patent laws attract more affiliates and also expand the size of existing affiliates in sectors with long product life cycles, with the largest effects in sectors with mid-length life cycles. This finding is important from a Southern welfare perspective, because it reveals that stronger patents expand not just the level but also the scope of local industrial activity. This latter effect strongly suggests that stronger patent laws attract greater levels of technology transfer in long life-cycle sectors. I also verify that responsiveness to patent laws is more pronounced among affiliates within low-productivity firms.

My article is related to several different literatures. The analysis contributes to a growing body of work that empirically evaluates the influence of intellectual property rights on foreign direct investment and technology transfer. Using data on direct investment in Eastern Europe and the former Soviet Union, Javorcik (2004) finds that stronger patent rights encourage firms to establish subsidiaries in high-technology sectors. Branstetter, Fisman, and Foley (2006) and Branstetter et al. (2011) examine firm-level responsiveness to patent reform events during the 1980s and 1990s; they find that the patent reforms are associated with increased intrafirm royalty payments and local affiliate activity among US multinationals, with the largest effects in high-patent firms. However, concurrent policy reforms and the high degree of correlation between measures of patent protection and general economic development pose a substantial challenge to empirical studies of this nature. I build on the foundation of these prior analyses by introducing systematic variation across sectors in product life-cycle lengths, a dimension that determines sensitivity to local patent laws but is theoretically independent of firms' sensitivity to general institutions and economic development. This variation enables me to isolate patent laws' causal influence on multinational activity, even in the presence of concurrent policy reforms.

My results complement prior studies that investigate differences across industries in the importance and effectiveness of patent protection. Firm-level surveys (Mansfield, Schwartz, and Wagner 1981; Levin et al. 1987; and Cohen, Nelson, and Walsh 2000) and other analyses (Schankerman 1998; Arora, Ceccagnoli, and Cohen 2003; and Moser 2005) reveal large differences in the usefulness of patents as a means of appropriating the returns from innovation, with patents conferring exceptionally effective protection in the chemical and medical equipment industries. I incorporate this insight into my analysis with an additional test and show that my results indeed hold with greater strength in sectors for which patents have been found to be effective. The results of my analysis also provide evidence that, under certain conditions, the rate of technological obsolescence itself forms an additional factor conditioning the relative importance of patent protection across industries.

The theoretical model in this article is closely related to an extensive literature on international product cycles, which has developed following Vernon (1966) and includes contributions by Krugman (1979), Helpman (1993), and Antràs (2005). These models evaluate the process by which the manufacture of products shifts from the North, where innovation occurs, to the South, where manufacturing costs are lower. Similar to the model in Antràs (2005), my model emphasizes the voluntary nature of firms' production location decisions, allowing relocation timing to be endogenously determined. My theoretical departure, relative to this prior literature, is the introduction of cross-industry variation in the economic durability of products and ideas.

This article also contributes to a line of research examining the impact of institutional frictions on foreign direct investment. Recent studies have emphasized the influence of financial development, investor protection laws, and contractual imperfections on multinational activity (Antràs, Desai, and Foley 2009; Manova, Wei, and Zhang 2011; Bernard et al. 2010; Antràs 2003). Others have suggested that the effects of these imperfections are acutely felt by innovative firms, particularly those seeking to manufacture cutting-edge technology abroad (Antràs and Helpman 2004; Nunn and Trefler 2008; Davidson and McFetridge 1985).

The rest of the article presents my theoretical and empirical analysis. In Section I, I develop a product cycle model with innovating firms that face imitation risk. After describing the data in Section II, I outline the estimation approach that will be used to evaluate the data in Section III. In Sections IV and V, I describe the empirical results. Section VI concludes.

I. Theory

The model developed below captures the essential trade-off between profit gains and offshore imitation risk faced by firms manufacturing abroad. The model formalizes the idea that this trade-off evolves over the product life cycle; production location decisions thus hinge on both host-country characteristics and product obsolescence rates. To highlight the theoretical results that correspond to my empirical analysis to follow, I present a simplified partial-equilibrium analysis with two countries. At the end of the section, I describe how my results extend to settings with multiple Southern countries and heterogeneous firms. Details appear in the online Appendix.

A. *Setup*

Time is continuous, and in each sector $j = 1, \ldots, J$ there is a continuum of horizontally differentiated product varieties. At every moment, firms in each industry manufacture varieties and sell to consumers in both the North and the South, symmetric countries with monopolistically competitive markets that are each of size 1. I assume innovating firms pay a one-time, sector-specific fixed cost to develop a new good, thereafter enjoying a monopoly in that good until it either is imitated or becomes obsolete. Sectors are distinguished by the pace of product (technology) obsolescence, which I assume is determined by technological developments specific to each industry but exogenous to individual firms.[5] Product life-cycle lengths T thus vary across sectors, but all sector-j products share the same life-cycle length T_j. This means that once a sector-j variety has reached a market maturity of T_j years, it becomes obsolete and is of no further economic value to consumers; intellectual property (both patented and tacit) associated with the retired variety also immediately becomes obsolete.[6]

To keep things simple, I treat the rate of new product entry as exogenous and assume that it is constant and equal to the rate of product obsolescence in each sector. This implies that obsolete varieties are immediately replaced by new innovations, leaving unchanged and exogenous the overall measure of varieties per sector, which I assume is 1 for all sectors j.[7] Notice that under my assumptions, the within-sector age distribution of product maturities at any point in time is uniform with density $\psi_j(t) \equiv 1/T_j$.

B. *Production*

Consider a firm that has introduced a particular good. To produce output, the firm uses its product-specific technology (proprietary information associated with the innovative good) to combine headquarters services and manufacturing. Innovating firms locate permanent headquarters in the North, and thus source headquarters services there, but at any time may costlessly shift manufacturing to an affiliate in the South. As a monopolist, a firm earns flow profit π^N when manufacturing in the North, and π^S when manufacturing in the South. I assume that economic conditions favor manufacturing in the South so that $\pi^S > \pi^N$; such conditions may include differences between North and South in regulatory costs, trade costs, and factor

[5] What is important for my empirical analysis is that the obsolescence rates faced by US firms do not respond significantly to changes in the patent laws of foreign countries. Taking obsolescence rates as technologically determined is a simple way to ensure this in my model, though it is a stronger assumption than I need for the empirical analysis. See Section IIIB for further discussion.

[6] This approach is consistent with the observation that most patented innovations are not renewed to the full term allowable by law (Schankerman and Pakes 1986), suggesting that the duration of a patent-based monopoly is, on average, a technological characteristic rather than a uniform legal standard. See Section IIA for further discussion.

[7] I discuss the implications of endogenous entry in later notes, but since my aim is to evaluate the production location decisions of innovating firms in response to the strength of foreign patent laws, this assumption seems reasonable provided that Northern multinationals make innovation decisions primarily on the basis of Northern patent laws.

prices.[8] In this setting, profit maximization implies a link between profit and revenue: each firm earns revenue $r^i = \sigma \pi^i$ when manufacturing in country i, where $\sigma > 1$ is the demand elasticity faced by all firms. Hence, firms earn higher revenues when manufacturing in the South, $r^S > r^N$.

An innovating firm's monopoly power may be disrupted by imitation, depending on its manufacturing location. Because manufacturing requires the use of product-specific proprietary information, imitators must obtain this information to enter. Consistent with evidence in Arora (1996), I assume that innovating firms protect proprietary information both formally (patents) and informally (trade secrets), and that all information—patented and tacit—is essential to successful production and must therefore be revealed to manufacturing employees.[9] By revealing trade secrets to these employees, I assume the firm enables potential imitators located near the firm's manufacturing facility to access this previously undisclosed information at lower cost.[10] In particular, I assume cross-border access to product-specific trade secrets is prohibitively costly, so that imitators are economically constrained to pursue only those varieties that have been locally manufactured.

Imitation risk affects innovators' location decisions, because entry by an imitator results in profit losses. Specifically, innovating firms competing with an imitator capture only a fraction of the per-period profits described above. This fraction depends on the quality of local patent enforcement, which I summarize with a pair of country-specific indexes ξ_N and ξ_S. Similar to Grossman and Lai (2004), ξ_i is the probability that a country-i patent will be enforced at any point in time but could be equivalently interpreted as the fraction of territory in which patents are enforced. I assume that patents are perfectly enforced in the North, but not in the South, $\xi_N = 1$ and $\xi_S < 1$. Only where a patent fails to be enforced may imitators' products compete directly with innovating firms'. Hence, imitation products may be sold only in the South.

C. *Imitation*

Potential imitators exist in the South. Any imitator with access to the proprietary information necessary for production may begin reverse-engineering a product; as

[8] Because the model provides a theoretical motivation for how international investment responds to patent reform events in the medium term, I do not include longer-run considerations such as endogenous changes in Southern profits π^S and Northern innovation rates. These general equilibrium effects are of clear theoretical and practical importance (Grossman and Helpman 1991) viewed from an aggregate perspective and will be accommodated by my main empirical specification that includes country-by-year fixed effects and, thus, emphasizes cross-sector rather than absolute comparisons in multinational activity.

[9] My assumption that firms protect product-specific proprietary technology both formally and informally is essential for the theoretical mechanism described here to be applicable; comprehensive patenting leaves imitators with little to learn once offshoring begins, while an absence of patents implies no reliance on patent institutions. That patented and tacit knowledge are assumed complements in the production process is also supported by firm-level evidence on knowledge bundling in Arora (1996). More general forms of tacit knowledge, not specific to a particular product, are also of clear importance (e.g., Klepper and Sleeper 2005); I consider firms' general knowledge implicitly in Sections IG and IVD.

[10] My assumption that acquiring proprietary information is less costly when it is in active local use is consistent with evidence in Poole (2013): employees of domestic establishments experience significant wage increases following hire of a former multinational employee, providing support for the existence of within-firm information transfers across workers. Relatedly, Moser (2005) finds complex innovations that are well protected by secrecy are significantly less likely to be patented. This supports the idea that firms' tacit intellectual property may be particularly difficult to discern by observing only final products and patents; even with access to patents and imported products from the North, imitators may therefore gain critical product information through direct access to manufacturing workers.

described above, this proprietary information is accessible in the South whenever manufacturing is located there, and it is otherwise not accessible to potential imitators. As in Grossman and Helpman (1991), Glass and Saggi (2002), and elsewhere, I assume the time to imitation success m is uncertain and that success arrives at a constant Poisson rate. For simplicity, I also assume the arrival time is restricted to a known interval $[0, \overline{m}]$, implying that m follows a uniform distribution over this period.[11] Imitation effort thus may or may not yield an imitation product within the targeted variety's lifetime.

A successful imitator competes with the innovating firm wherever patents are not enforced (Grossman and Helpman 1991) until the variety becomes obsolete. I assume that imitators are able to profitably capture any such market where patents are not enforced. Because imitation is costless, this implies that offshoring manufacturing immediately raises the hazard of imitation.

D. *The Product Cycle*

Firms make profit-maximizing location decisions, taking into account the behavior of Southern imitators. Specifically, each sector-j firm selects the optimal product maturity $t_j^* \in [0, T_j]$ at which to begin offshoring production in the South by maximizing lifetime expected profits,

$$\text{(1)} \qquad \mathrm{E}_m[\Pi_j(t)] = 2\pi^N t + 2\pi^S \mathrm{E}_m[\min\{T_j - t, m\}] + (1 + \xi_S)\pi^S \mathrm{E}_m[\max\{0, T_j - t - m\}].$$

Equation (1) reveals the effect of imperfect patent enforcement in the South on Northern firms' offshoring incentives. Selling to both markets, the firm earns $2\pi^N$ with certainty until offshoring begins at maturity t, but thereafter faces imitation risk. The firm then earns full profit $2\pi^S$ for the length of time m if imitation precedes obsolescence, and for $T_j - t$ otherwise. In the former case, once imitation has occurred, profit is π^S only where patents are enforced, namely in the Northern market and in a fraction ξ_S of the Southern market; postimitation profit is thus $(1 + \xi_S)\pi^S$.

Maximization of $\mathrm{E}_m[\Pi_j(t)]$ over possible relocation maturities t reveals a time-to-obsolescence offshoring cutoff $\tau^*(\xi_S)$ that is invariant across industries with different product life-cycle lengths T_j. Varieties with less than $\tau^*(\xi_S)$ time remaining before obsolescence are manufactured in the South, while all other varieties are manufactured in the North. The value of this cutoff,

$$\text{(2)} \qquad \tau^*(\xi_S) = \frac{\pi^S - \pi^N}{\pi^S(1 - \xi_S)} 2\overline{m},$$

depends on ξ_S, innovators' π^N and π^S, and imitators' $\overline{m}$. Notice that firms manufacture in the South earlier in the product life cycle when Southern patents are better protected: $\partial\tau^*(\xi_S)/\partial\xi_S > 0$.

[11] To reduce the taxonomy of cases, I also assume that $\overline{m} > \max_j\{T_j\}$, so that the time to imitation success is relatively uncertain compared with the length of product life cycles.

LEMMA 1: *In industries with $T_j < \tau^*(\xi_S)$, products are always manufactured in the South. In all other industries, product cycles emerge: manufacturing is initially in the North but relocates to the South once the product has $\tau^*(\xi_S)$ time remaining before obsolescence and remains there until obsolescence. The cutoff $\tau^*(\xi_S)$ is increasing in Southern patent protection ξ_S: $\partial\tau^*(\xi_S)/\partial\xi_S > 0$.*

Lemma 1 implies that products are manufactured in the North for $t_j^*(\xi_S) \equiv \max\{0, T_j - \tau^*(\xi_S)\}$ time, and in the South for $T_j - t_j^*(\xi_S) = \min\{T_j, \tau^*(\xi_S)\}$ time. This further implies that the measure of products manufactured in the South at any point in time is

$$(3) \qquad S_j(\xi_S) \equiv \int_{t_j^*(\xi_S)}^{T_j} \psi_j(t)\, dt = \frac{T_j - t_j^*(\xi_S)}{T_j} = \min\left\{1, \frac{\tau^*(\xi_S)}{T_j}\right\},$$

which is weakly decreasing in T_j; this reflects the essential point that longer-lived products spend (weakly) more of their lives being produced in the North, and, hence, at any point in time there are (weakly) fewer of them manufactured in the South. I summarize this as follows:

RESULT 1: *The measure of varieties manufactured in the South in sector j is (weakly) decreasing in product life-cycle lengths T_j: $\partial S_j(\xi_S)/\partial T_j \leq 0$.*

Notice that $S_j(\xi_S)$ may also be interpreted as the fraction of offshored varieties in sector j. In the top panel of Figure 1, I plot $S_j(\xi_S)$ (vertical axis) as a function of sectors' product life-cycle lengths T_j (horizontal axis) where T_j ranges between $T_{min} < \tau^*(\xi_S) < T_{max}$; I show $S_j(\xi_S)$ at two different levels of Southern patent protection, ξ_S (solid) and ξ_S' (dashed), with $\xi_S < \xi_S'$. Furthermore, Southern affiliates in sector j earn aggregate expected revenues $R_j(\xi_S)$ as follows, obtained by integrating variety-specific revenues over the distribution of product maturities:

$$(4) \quad R_j(\xi_S) = \int_{t_j^*(\xi_S)}^{T_j} \left(2r^S(1 - \kappa_{im}(t)) + (1+\xi_S)r^S\kappa_{im}(t)\right)\psi_j(t)\, dt$$

$$= \begin{cases} 2r^S \cdot \left(1 - \dfrac{T_j}{2\overline{m}}\right) + (1+\xi_S)\cdot r^S \cdot \dfrac{T_j}{2\overline{m}}, & T_j < \tau^*(\xi_S) \\[2ex] 2r^S \cdot \left(\dfrac{\tau^*(\xi_S)}{T_j} - \dfrac{\tau^*(\xi_S)^2}{2\overline{m}\,T_j}\right) + (1+\xi_S)\cdot r^S \cdot \dfrac{\tau^*(\xi_S)^2}{2\overline{m}\,T_j}, & T_j \geq \tau^*(\xi_S), \end{cases}$$

where $\kappa_{im}(t)$ is the probability that a maturity-t product is currently imitated. For the first case in (4), products are always manufactured in the South but face imitation risk; at any moment, a fraction $\frac{T_j}{2\overline{m}}$ have been imitated and thus earn only $(1+\xi_S)\, r^S$. In the second case, products have longer lifetimes and only a measure $\frac{\tau^*(\xi_S)}{T_j}$ are manufactured in the South at any time; of these, a smaller measure $\frac{\tau^*(\xi_S)^2}{2\overline{m}\,T_j}$ have been imitated and earn only $(1+\xi_S)r^S$.

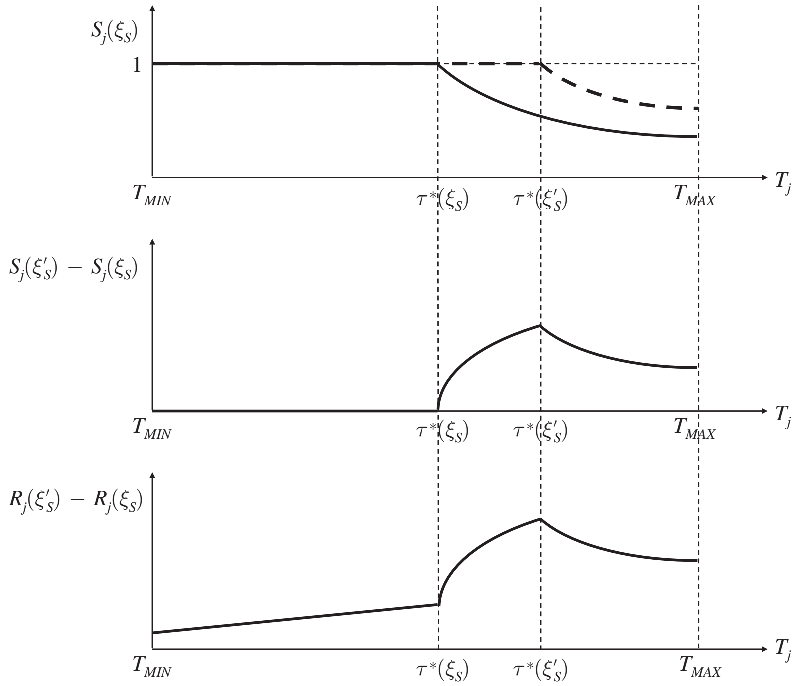

Figure 1. Product Life-Cycle Lengths and the Location of Manufacturing Activity

Notes: The upper panel of this figure shows the measure of product varieties manufactured in the South $S_j(\xi_S)$ as a function of product life-cycle lengths T_j at two different levels of host-country patent protection ξ_S (solid) and ξ'_S (dashed), $\xi_S < \xi'_S$. The middle panel takes the difference between these two curves, showing that countries with stronger patent protection attract manufacturing activity for a larger measure of varieties, but only in sectors with $T_j > \tau^*(\xi_S)$. The lower panel shows the difference in affiliate revenues at the two levels of patent protection, which increase due to the combined effect of entry and stronger protection of existing imitated varieties.

E. *Cross-Sector Response to Strengthened Patent Rights*

Suppose the South enacts a policy change that strengthens local patent enforcement from ξ_S to $\xi'_S > \xi_S$. By equation (2), firms optimally offshore manufacturing earlier in the product life cycle following this change. From (3), a straightforward implication of this is that the difference in the measure of varieties manufactured in the South at ξ_S and at $\xi'_S > \xi_S$ depends on T_j as follows:

$$S_j(\xi'_S) - S_j(\xi_S) = \begin{cases} 0, & T_j < \tau^*(\xi_S) \\ \dfrac{T_j - \tau^*(\xi_S)}{T_j}, & T_j \in [\tau^*(\xi_S), \tau^*(\xi'_S)] \\ \dfrac{\tau^*(\xi'_S) - \tau^*(\xi_S)}{T_j}, & T_j > \tau^*(\xi'_S) \end{cases}. \quad (5)$$

I plot $S_j(\xi'_S) - S_j(\xi_S)$ in the middle panel of Figure 1. This extensive-margin effect can be interpreted as the measure (level) of varieties that is immediately offshored to the South following the patent reform. This effect may also be interpreted as a change in the fraction of offshored varieties in sector j; my empirical approach below is consistent with either interpretation. In industries with product life cycles shorter than the original sourcing cutoff $(T_j \le \tau^*(\xi_S))$, firms do not respond to the reform; in these industries, varieties were already manufactured in the South for their full lifetime at ξ_S and continue to be at ξ_S'. In industries with longer product life cycles $T_j > \tau^*(\xi_S)$, however, firms shift the manufacture of marginal varieties to the South. Marginal varieties are the subset of products with between $\tau^*(\xi_S)$ and $\tau^*(\xi_S')$ remaining years before obsolescence and are, thus, found only in industries with $T_j > \tau^*(\xi_S)$. The measure of marginal varieties is increasing in T_j for $T_j \in [\tau^*(\xi_S), \tau^*(\xi_S')]$, and is decreasing in T_j for $T_j > \tau^*(\xi_S')$. $S_j(\xi_S') - S_j(\xi_S)$ is thus a nonmonotonic function of T_j.

RESULT 2: *The increase in the measure of sector-j varieties manufactured in the South $S_j(\xi_S') - S_j(\xi_S)$ following a patent reform from ξ_S to ξ_S' is a nonmonotonic function of T_j. Specifically, it is zero for $T_j < \tau^*(\xi_S)$, increasing for $T_j \in [\tau^*(\xi_S), \tau^*(\xi_S')]$, and decreasing for $T_j > \tau^*(\xi_S')$. The largest impact occurs in the industry with $T_j = \tau^*(\xi_S')$.*

Revenues earned by Southern affiliates are impacted by patent reform both because of newly shifted manufacturing and because existing imitated varieties capture a larger share of Southern sales under the stronger patent regime. Building from equations (4) and (5), the change in industry-j revenues earned offshore following a Southern patent reform from ξ_S to ξ_S' depends on T_j

$$(6)\quad R_j(\xi_S') - R_j(\xi_S) = \begin{cases} r^S \cdot (\xi_S' - \xi_S) \cdot \dfrac{T_j}{2\overline{m}}, & T_j < \tau^*(\xi_S) \\[2ex] r^S \cdot \left(2\,\dfrac{T_j - \tau^*(\xi_S)}{T_j} + (\xi_S' - \xi_S) \cdot \dfrac{[\tau^*(\xi_S)]^2}{2\overline{m}\,T_j} \right), & T_j \in [\tau^*(\xi_S), \tau^*(\xi_S')] \\[2ex] r^S \cdot \left(2\,\dfrac{\tau^*(\xi_S') - \tau^*(\xi_S)}{T_j} + (\xi_S' - \xi_S) \cdot \dfrac{[\tau^*(\xi_S)]^2}{2\overline{m}\,T_j} \right), & T_j > \tau^*(\xi_S'). \end{cases}$$

$R_j(\xi_S') - R_j(\xi_S)$ above is also a nonmonotonic function of T_j, increasing for $T_j \le \tau^*(\xi_S')$ and decreasing for $T_j > \tau^*(\xi_S')$, implying that the largest response to the patent reform is in the industry with $T_j = \tau^*(\xi_S')$. I illustrate $R_j(\xi_S') - R_j(\xi_S)$ as a function of T_j in the bottom panel of Figure 1.

RESULT 3: *Following a patent reform from ξ_S to ξ_S', the increase in sector-j revenues earned by Southern affiliates $R_j(\xi_S') - R_j(\xi_S)$ is a nonmonotonic function of T_j. Specifically, it is increasing for $T_j \le \tau^*(\xi_S')$, and decreasing for $T_j > \tau^*(\xi_S')$. The largest impact occurs in the industry with $T_j = \tau^*(\xi_S')$.*

Results 2 and 3 form the foundation for my empirical strategy (Section III).

F. *Multiple Southern Countries*

Qualitatively identical results to those described above apply to comparisons in the cross-section between countries with different levels of patent protection. Consider two Southern countries S and S' that are identical but for different patent institutions ξ_S and ξ_S', with $\xi_S' > \xi_S$, and for simplicity, assume trade costs are negligible.[12] Firms' offshoring incentives are captured by an equation similar to (1) that incorporates additional considerations. First, firms optimally consolidate manufacturing activity in a single location that serves as an export platform for all other markets (e.g., Arkolakis et al. 2013). Second, imitators in the South may sell both locally and in other markets where patents are imperfectly enforced; however, because exporting imitation goods requires evading both local and foreign patent authorities, the export market enjoyed by imitators is small compared with the local market. Finally, because multinational activity is observed in countries with both strong and weak patent protection, in this multicountry extension I allow firms' affiliate location decisions to depend both on maximized lifetime profits as a function of patent protection (1) and on an unobserved, location-specific profit component. I assume this additional component enters the lifetime profit function and is drawn independently across country-variety pairs according to a known distribution.

Results 2 and 3 apply to Southern countries S and S' collectively in this setting. In particular, it can be shown (see the online Appendix) that, similar to (5), the industry-level measure of varieties manufactured only in the stronger-patent host country $S_j(\xi_S') - S_j(\xi_S)$ is a nonmonotonic function of T_j. And, comparing total affiliate revenues in the stronger-patent host country $R_j(\xi_S')$ with affiliate revenues in the weaker-patent host country $R_j(\xi_S)$ reveals a difference analogous to (6) and similarly nonmonotonic in T_j.

RESULT 4: *The differences* (i) *between the measure of sector-j varieties manufactured in the host country with stronger patent protection* ξ_S' *versus in the host country with weaker patent protection* ξ_S, $S_j(\xi_S') - S_j(\xi_S)$, *and* (ii) *between the revenues earned by sector-j Southern affiliates in the host country with stronger patent protection* ξ_S' *versus in the host country with weaker patent protection* ξ_S, $R_j(\xi_S') - R_j(\xi_S)$, *are both nonmonotonic functions of* T_j. *Specifically, both differences are constant for* $T_j < \tau^*(\xi_S)$, *increasing for* $T_j \in [\tau^*(\xi_S), \tau^*(\xi_S')]$, *and decreasing for* $T_j > \tau^*(\xi_S')$. *Both* $S_j(\xi_S') - S_j(\xi_S)$ *and* $R_j(\xi_S') - R_j(\xi_S)$ *are therefore largest in the industry with* $T_j = \tau^*(\xi_S')$.

G. *Firm Heterogeneity*

In the presence of firm heterogeneity in productivity, patent laws have differential effects on production location decisions across firms with different productivity levels. Returning to the two-country North-South world, I assume that within-sector productivity differences across innovators can be summarized by a

[12] Results in this section also obtain in settings with higher transport costs; see the online Appendix.

positive firm-specific parameter $\varphi \in [\varphi_L, \varphi_H]$ that affects profits as in Melitz (2003). Firms with higher productivity draws earn higher profits whether manufacturing in the North $\pi^N(\varphi) = \pi^N \varphi^{\sigma-1}$ or in the South $\pi^S(\varphi) = \pi^S \varphi^{\sigma-1}$, where π^N, π^S, and $\sigma > 1$ are as defined above. Assume further that Southern firms share a fixed productivity level below φ_H, and that imitators compete with innovating firms on the basis of price wherever patents are not protected.[13] This monopolistically competitive setting implies a link between profits, prices, and marginal production costs, further implying that the Northern innovator may be productive enough that successful imitators are unable to capture unprotected markets. Specifically, it can be shown (see the online Appendix) that firms with sufficiently high productivity levels φ are unaffected by imitation in the South, and are therefore insensitive to the quality of Southern patent protection ξ_S.

RESULT 5: *The differential effects of improved Southern patent protection are more pronounced for relatively unproductive firms, because high-productivity firms are not affected by imitation in the South and are therefore not responsive to patent reform.*

II. Data and Descriptive Statistics

Evaluating the theoretical results above requires measures of product life-cycle lengths, multinational activity, and host-country intellectual property rights. I describe these below.

A. *Product Life-Cycle Lengths by Industry*

One natural approach to measuring product life-cycle lengths T_j is to determine the average retail duration of individual product varieties. Direct measures of product turnover are available in some industries (for example, Broda and Weinstein 2010), and I will describe results based on such an index later in Section V. But the theory developed in Section I suggests an alternative measurement approach. In the model, each product is associated with proprietary information (patented and tacit intellectual property) that is necessary for production, but which loses its economic value when the product becomes obsolete. Implicit in this view is a separability of proprietary information across products: the information required to manufacture each new product must be distinct from that required for any older product.[14] The critical difference across industries in the model is therefore not the time lapse between different versions of a product embodying the same innovative idea, but

[13] My assumption that Southern firms do not inherit innovators' productivity through imitation is consistent with evidence in Qian (2008), Bloom and Van Reenen (2007), Bloom, Sadun, and Van Reenen (2012), and Bloom et al. (2013), that firm productivity is composed of nonproduct characteristics including manufacturing ability, management quality, corporate structure, marketing, sales networks, and so on. These elements are difficult to imitate but essential to maintaining productivity over time (Teece, Pisano, and Shuen 1997); for example, Qian (2008) provides detailed evidence that imitators are unable to inherit innovators' manufacturing ability in the footwear sector. Notice that my results require only that productivity spillovers between innovators and imitators are incomplete.

[14] This is important because, theoretically, a firm's manufacturing location impacts its imitation risk only if products are technologically distinguished enough to preclude imitation of one product solely on the basis of an ability to imitate another.

rather the economic lifetime of the innovative idea itself, which may span several successive versions over the product life cycle as in Klepper (1996).

As an illustration of this product definition, consider the example of automobiles. New car models within an automobile product line are introduced annually (termed the "model cycle" in Bils 2009), but the technological overlap across successive models is substantial.[15] Because of this technological overlap, assigning T_j to be one year for automobiles would be incorrect in the context of the theory in Section I; firms facing imitation lose exclusive control of proprietary technologies, impacting the profitability of both current and future models. From an imitation view, the relevant product life cycle for cars is, thus, longer than a year, and according to the model in Section I, coincides with the economic lifetime of embedded technologies.[16]

Motivated by this product definition, my measurement approach aims to capture cross-industry variation in product life-cycle lengths by determining the economic durability of embedded technologies. To do this, I use detailed information in US patent records. The length of time during which a given patent continues to be cited by subsequent patents provides an indication of the market lifetime of the cited patent's technology, with a long average "forward citation lag"—the time lapse between the cited patent's grant date and a subsequent citation—indicating that the technology exhibits lasting relevance to future innovation.[17] In the online Appendix, I formalize my measurement approach within a simple theory of patent citations and derive a sector-level index of product life-cycle lengths that may be constructed using standard datasets (NBER US Patent Citation Data File; Hall, Jaffe, and Trajtenberg 2001). The index, $\hat{T}_j$, is the average forward citation lag within an industry.[18] To construct $\hat{T}_j$, I place no restriction on the patent class of citing patents, to allow for the possibility that a patented technology may have relevance to future innovation not only within its own class, but also for other patent classes. Notice that the validity of this $\hat{T}_j$ rests on the assumption that the durability of patented product-specific knowledge is positively correlated with the durability of tacit (unpatented) product-specific knowledge; this is consistent with the idea that both are essential elements of the same final good.[19]

[15] Successive versions of the Honda Accord, for example, are so similar that the BLS substitutes new versions for old (e.g., the 2012 Honda Accord LX is substituted for the 2011 Honda Accord LX, with minimal adjustment) to establish price comparisons underlying official US inflation indexes (Bils 2009).

[16] See also Klepper (1996) and Greenstein and Wade (1998) for further discussion and references.

[17] As an example of the functional role played by patent citations, consider the medical device firm Macquet Cardiovascular. According to an executive at the firm, patents for follow-on inventions (including theirs) are compelled to cite prior art as an indication of close or overlapping claims that may result in a royalty payment. Such citations, and royalty payments, continue to accrue only as long as the prior technology is viable and cease when it becomes obsolete. This logic is related to Narin and Olivastro (1993) and Mehta, Rysman, and Simcoe (2010), and under general conditions implies a link between citation timing and technology obsolescence at the industry level (see online Appendix A.2).

[18] As an illustration, consider two US patents with application dates in 1991. US patent 5,140,955 granted in 1992 to Honda Motor Company (class 123, internal combustion engines) was cited 29 times (primarily by GM and Caterpillar) during its first 15 years, and the average time between its grant date and a subsequent citation is seven years, with the most recent citation in 2005. By comparison, US patent 5,239,631 granted in 1993 to IBM Corporation (class 710, electrical computers and digital data processing systems) was cited a comparable 25 times (primarily by Intel and Compaq) during its first 15 years; however, its average forward citation lag is much shorter, 3.6 years, and the most recent citation is in 2001. In line with this example, the constructed product life-cycle length index is longer for motor vehicles (SIC 371) at 9.64 years than for computers and office equipment (SIC 357) at 8.38 years.

[19] This assumption is also consistent with evidence indicating complementarity between patented and tacit knowledge (Arora 1996). In addition, it is important that most citations are inserted by official patent examiners, making strategic citations improbable (e.g., Sampat 2010).

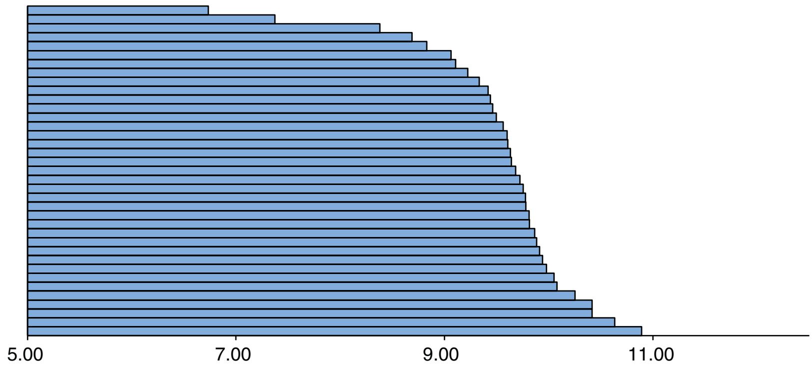

Figure 2. Average Citation Lag (*Years*)

Notes: This figure shows the average patent citation lag for each of the 37 SIC three-digit industries studied in this paper. Citation lags are measured in years and were computed using data in the NBER Patent Citation Datafile (Hall, Jaffe, and Trajtenberg 2001).

Mapping $\hat{T}_j$ onto data for multinational activity from the BEA involves using a concordance between US patent classes and SIC three-digit industry codes.[20] In most cases, many patent classes correspond to the same SIC industry. I therefore take an unweighted average $\hat{T}_j$ index across all patent classes matched to a given SIC code.

There is considerable variation in this index across industries as shown in Figure 2. In Table 1, I provide the names of the top five and bottom five industries, as well as the five industries around the median, ranked by measured product life-cycle lengths. Electronics, clockwork-operated devices, and computers have some of the shortest product life cycles, while machine products, shipping containers, and nonelectric heating equipment have some of the longest life cycles. Table 1 also indicates that my index of product life-cycle lengths does not simply reflect sectors' innovation intensities. While some sectors with quickly obsolescing technologies (low T_j) have higher R&D intensities (the ratio of R&D to sales), a plot of R&D intensity against T_j in Figure 3 shows that these two measures reflect distinct industry characteristics (correlation -0.32). My empirical analysis explicitly incorporates sectors' R&D intensities and other characteristics that may affect firms' sensitivity to patent protection (Sections III–V). I document the stability of $\hat{T}_j$ over time, in both magnitude and the rank-ordering across industries, in the online Appendix.

B. *Comparing $\hat{T}_j$ and Alternative Measures of Technology Obsolescence*

For robustness, in Section V I use several alternative proxies for T_j. First, skewness in patent values (Schankerman 1998; Hall, Jaffe, and Trajtenberg 2001) suggests that patents at the upper tail of each distribution are relatively successful innovations that, as a result, may better reflect the durability of technology multinational firms are concerned about protecting. Accordingly, I construct alternative proxies

[20] This concordance can be downloaded from the website of the US Patent and Trademark Office, http://www.uspto.gov/web/offices/ac/ido/oeip/taf/data/sic_conc/2005_diskette/.

Table 1—Product Life-Cycle Lengths (T) and R&D Intensities by Sector

Industry name	Rank	SIC Code	T (years)	R&D intensity
Shortest product life cycles				
Electronics machinery	1	383	6.73	0.0527
Watches, clocks, clockwork operated devices	2	387	7.37	0.0239
Computer and office equipment	3	357	8.38	0.0987
Agricultural chemicals	4	287	8.69	0.0219
Electronic components and accessories	5	367	8.83	0.242
Intermediate-length product life cycles				
Miscellaneous industrial and commercial	19	359	9.68	0.026
Miscellaneous chemical products	20	289	9.73	0.029
Surgical, medical, dental instruments, and supplies	21	384	9.75	0.058
Farm and garden machinery and equipment	22	352	9.78	0.020
Household appliances	23	363	9.78	0.012
Longest product life cycles				
Fabricated structural metal products	33	344	10.25	0.0102
Cutlery, handtools, and general hardware	34	342	10.41	0.0137
Screw machine products, bolts, nuts, screws	35	345	10.42	0.0240
Metal cans and shipping containers	36	341	10.63	0.0119
Heating equipment, except electric	37	343	10.89	0.00986

Notes: This table shows the product life-cycle length (the average patent citation lag) and R&D intensity (the ratio of R&D to sales) for the top and bottom five industries ranked by product life-cycle lengths, as well as five industries around the median. Patent citation data is from the NBER US Patent Citation Data File (Hall, Jaffe, and Trajtenberg 2001). R&D intensity is the average ratio of R&D to sales by industry among multinationals and is based on firm-level data from the Bureau of Economic Analysis.

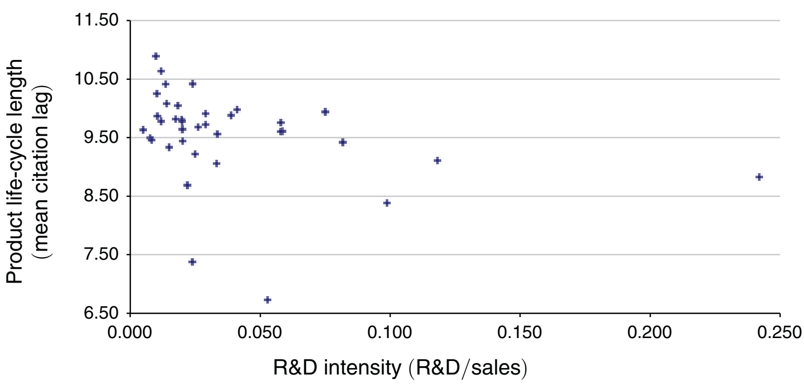

Figure 3. Product Life-Cycle Lengths versus R&D Intensities across Industries

Notes: This figure plots the product life-cycle length against the R&D intensity of each SIC three-digit industry studied in this paper. Product life-cycle lengths are average patent citation lags and are computed using data in the NBER Patent Citation Datafile. R&D intensity is the R&D-to-sales ratio in the BEA firm-level data during the benchmark years between 1982 and 2004.

for T based on the seventy-fifth and eighty-fifth percentiles of each citation lag distribution. Second, I consider measures of knowledge obsolescence and patent duration that are based on patent renewal data. Prior research has used patent renewal data to evaluate patent quality, the rate of depreciation of knowledge capital, and, relatedly, the asset value of patents (Pakes 1986; Schankerman and Pakes 1986;

Lanjouw, Pakes, and Putnam 1998; Lanjouw and Schankerman 2004; Schankerman 1998; Bessen 2008; Mehta, Rysman, and Simcoe 2010). I construct a disaggregated renewal-based measure and use this as an alternative proxy for T.[21] Finally, I compare my main results with specifications using the Broda and Weinstein (2010) product turnover index, as described above. Further detail on data construction and the comparison across alternative measures of T_j may be found in the online Appendix.

C. *US Multinational Activity Abroad*

I use firm-level panel data on the global operations of US-based multinationals from the Bureau of Economic Analysis (BEA) Survey of US Direct Investment Abroad. These data provide information on US parent companies and each foreign affiliate on an annual basis.[22] This analysis uses data from benchmark-year surveys, which are the most extensive in both scope and coverage and are available for 1982, 1989, 1994, 1999, and 2004.[23] Table 2 provides a summary of multinational activity during the five benchmark years by industry, including total assets, employment, sales, and R&D expenditures across the countries and industries in this study.[24]

To analyze the influence of patent laws on multinationals' affiliate location decisions, I use disaggregated information on the sales, employment, and assets of multinational affiliates located in 92 countries. Countries were included in the dataset if (i) any US FDI was recorded in any of the benchmark years, and (ii) the patent rights index described below was available for the host country in at least two periods. Affiliate sales are reported by three-digit industry code, making it possible to categorize affiliate activity by primary industry. I also compare affiliate sales with US exports by sector using data from the US Census Bureau.[25]

[21] Direct comparisons between existing knowledge-depreciation rates (e.g., Schankerman 1998) and $\hat{T}_j$ are difficult, because these measures are available for broad technology categories, while $\hat{T}_j$ is available for hundreds of patent classes. I therefore construct a more disaggregated renewal-based measure and find a correlation with $\hat{T}_j$ of approximately 0.30. This suggests that citation lag and renewal data, though distinct, may capture overlapping information regarding the durability of patented technology.

[22] Any US person having direct or indirect ownership or control of 10 percent or more of the voting securities of an incorporated foreign business enterprise or an equivalent interest in an unincorporated foreign business enterprise at any time during the benchmark fiscal year in question is considered to have a foreign affiliate. However, for very small affiliates that do not own another affiliate, parents are exempt from reporting with the standard survey form. Foreign affiliates are required to report separately unless they are in both in the same country and three-digit industry. Each affiliate is considered to be incorporated where its physical assets are located.

[23] The BEA's data coverage is nearly complete: in a typical benchmark year, the survey accounts for over 99 percent of affiliate activity. In 1994, for example, participating affiliates accounted for an estimated 99.8 percent of total assets, 99.7 percent of total sales, and 99.9 percent of total US FDI. This reflects the requirement of participation for every US person having a foreign affiliate. However, in certain cases involving missing survey responses, the BEA data may instead report imputed values; these values are coded accordingly, and I exclude from my analysis all such observations.

[24] While these data provide a nearly complete representation of offshore manufacturing by US firms by country, a limitation is that coverage is restricted to activity occurring within firm boundaries. As a result, my estimates may provide only a lower bound for the overall effect of patent protection on industrial activity, as stronger patent protection may attract not only increased within-firm manufacturing, but also greater levels of arms' length sourcing. Both modes of response are consistent with my theoretical model.

[25] Sector-level trade data may be obtained directly from the Census Bureau or may alternatively be downloaded from Peter Schott's website, http://www.som.yale.edu/faculty/pks4/.

Table 2—Regression Summary Statistics

Variable	Mean	SD	Min	Max
Country-industry-year level				
Total affiliate sales (thousands $US)	257,000	2,000,000	−9,100	88,600,000
Total affiliate employment	942	5,300	0	161,000
Total affiliate assets (thousands $US)	316,000	3,120,000	0	171,000,000
Total affiliate R&D (thousands $US)	3,300	42,800	0	2,131,732
Number of affiliates	2.31	7.48	0	114
Number of affiliates per country-year, conditional on presence	7.98	12.2	1	114
Affiliate-year level				
Affiliate sales (thousands $US)	80,200	491,000		
Affiliate employment	310	1,290		
Affiliate assets (thousands $US)	79,500	439,000		
Affiliate R&D (thousands $US)	1,070	16,400		
Affiliate sales to local unaffiliated (only majority owned)	43,000	210,000		
Affiliate sales to US (only majority owned)	10,900	214,000		
Industry level				
Average patent citation lag (years)	9.55	0.79	6.73	10.9
Average eighty-fifth percentile patent citation lag	17	1.18	13	20
Average R&D intensity	0.0379	0.0435	0.00502	0.242
Average share of intrafirm sales $S/(S+X)$	0.22	0.36	−0.107	1.4
log US exports (thousands $US)	74,276.60	618,355.30	0	3.90E + 07
Country-year level				
Patent index	2.76	1.10	0.59	4.67
Delta patent index	0.32	0.47	0	2.18
log GDP per capita	8.79	1.11	5.08	11.1
Corporate tax rate	0.09	0.07	0.0019	0.89
General				
Number of parent companies per year	959	96	886	1,125
Number of affiliates per parent	9.74	20.5		
Number of countries	94			
Number of industries	37			
Number of observations, industry level	13,629			
Number of observations, industry level, conditional on presence	4,783			
Number of observations, affiliate level	22,505			

Notes: This table summarizes multinational activity, host-country institutions, and industry characteristics across 94 countries, 37 industries, and all benchmark years during 1982–2004. Average patent citation lags were calculated using the NBER US Patent Citation Data File (Hall, Jaffe, and Trajtenberg 2001) by patent class and matched to three-digit SIC industry codes using a standard USPTO concordance. GDP per capita is from the Penn World Table (Heston, Summers, and Aten 2009). All other variables are from the Bureau of Economic Analysis Survey of US Direct Investment Abroad, and pertain to US outward foreign direct investment during the five most recent benchmark survey years (1982, 1989, 1994, 1999, 2004).

D. *Intellectual Property Rights Protection across Countries*

A proxy for the strength of patent protection across countries is provided by an index developed in Ginarte and Park (1997) and updated in Park (2008). This index is widely used because of its detailed construction and extensive coverage.[26] The index documents the strength of patent rights in five distinct categories: (i) extent of coverage, (ii) membership in international patent agreements, (iii) provisions for loss of patent protection, (iv) enforcement mechanisms, and (v) duration of protection. Each category is given a score between zero and one based on whether

[26] See, for example, Qian (2007); Branstetter, Fisman, and Foley (2006); Javorcik (2004); McCalman (2004); and Yang and Maskus (2001).

prevailing patent laws meet specific, objective criteria.[27] The overall index is the unweighted sum of these five subindexes and, thus, ranges between zero and five, with higher values indicating stronger protection. My main results are based on this overall index; I also evaluate the importance of each individual subindex in alternative specifications. These indexes are available for 122 countries during 1980–2005 in five-year intervals; I match the year of each benchmark survey to the closest available index year. Summary statistics appear in Table 2.

III. Econometric Framework

The model in Section I has implications for the distribution of multinational activity across countries with different levels of intellectual property protection, and across industries with different product life-cycle lengths. In this section, I describe the estimation approach used to evaluate the data.

A. *Baseline Estimating Equation*

Results 2–4 state that the sensitivity of multinationals' offshored manufacturing activity to host-country patent protection is a nonmonotonic function of sectors' product life-cycle lengths (Figure 1). These results motivate an estimating equation of the following form:

$$(7)\quad \ln\left(MNC_{ijt}\right) = \beta + \gamma_1 \cdot IPR_{it} \times T_j + \gamma_2 \cdot IPR_{it} \times T_j^2 + \eta_{it} + \eta_j + \epsilon_{ijt},$$

where MNC_{ijt} is a measure of multinational activity in country i and industry j during year t, IPR_{it} is the patent protection index in country i and year t, and T_j represents the product life-cycle length of sector j. The main coefficients of interest γ_1 and γ_2 jointly capture the differential influence of patent laws on multinational activity across sectors with different T_j values. To be consistent with the model, these coefficients must satisfy several criteria: $\gamma_1 > 0$, $\gamma_2 < 0$, and $\gamma_1 \cdot T_j + \gamma_2 \cdot T_j^2 \geq 0$ across the observed range of T_j (6.7 to 10.9 years), reaching a peak within this range: $T_{peak} = \frac{-\gamma_1}{2\gamma_2} \in [6.7,\ 10.9]$.[28]

I estimate baseline specification (7) with several measures of affiliate activity MNC_{ijt}. I first estimate (7) using an indicator for positive affiliate sales (in country i, industry j, and period t) as the dependent variable. I then evaluate implications of Result 4 concerning the measure of offshored varieties by defining MNC_{ijt} to be the number of US firms with a foreign affiliate in country i, industry j, and period t. Implications for offshore revenues are evaluated by defining MNC_{ijt} to be affiliate

[27] For example, the enforcement mechanisms category was scored by adding binary indicators corresponding to the availability of (i) preliminary injunctions, (ii) contributory infringement pleadings, and (iii) burden-of-proof reversals. A country with laws meeting all three criteria would receive a value of 1 for this category.

[28] As described in Results 2–4, the peak impact of a Southern patent reform (or the cross-section analog) occurs in the sector with product life-cycle lengths equal to the new sourcing cutoff $\tau^*(\xi_S')$. This cutoff is implicitly constrained in (3) to be less than T_{MAX} and cannot fall below T_{MIN} without contradicting the model's result that products with short life cycles are relatively insensitive to host-country patent laws.

sales revenues earned in country i, industry j, and period t; I also use affiliates' assets and employment as alternative dependent variables.[29]

To evaluate Result 5 concerning the influence of firm-level productivity on changes in the size of individual affiliates, I estimate an affiliate-level version of (7) using sales, assets, and employment by affiliate. Notice that if US multinationals are multiproduct firms, Results 4 and 5 both have implications for affiliate-level expansion. Specifically, Result 4 indicates that affiliate size will respond to patent protection nonmonotonically as a function of product life-cycle lengths, while Result 5 states that this expansion will be concentrated within less productive firms.

The baseline specification above includes a number of important controls. Industry fixed effects η_j absorb omitted sector characteristics, including imitation timing $\overline{m}$, per-period profit π^N and π^S, factor intensities of production, the average productivity of firms, total industry size, preferences, and the main effect of T_j. Country-year fixed effects η_{it} control for country characteristics that affect multinational activity and may (or may not) change over time, including development levels, size, factor costs, corporate tax rates, competition levels, and also IPR_{it}. The error term ϵ_{ijt} combines any omitted factors that affect firms' offshore production patterns. Because there may be measurement error in the index of host-country patent protection, I cluster errors by country in all reported results, but the results are robust to alternative levels of clustering.

B. *Identification*

Identification of γ_1 and γ_2 is based on within–"country-year" variation across industries in product life-cycle lengths. A key advantage of this approach is that it mitigates the empirical challenge introduced by concurrent policy reforms. Improvements to intellectual property rights often occur at the same time as more general institutional or economic changes that influence the location of multinational activity.[30] Upon joining the WTO, for example, countries align patent protection with the standards in the Agreement on Trade Related Aspects of Intellectual Property Rights (TRIPS), but also make substantial changes to domestic policies, including trade and investment barriers. Because such factors affect multinationals' investment decisions, the interpretation of coefficients in a simple regression of MNC_{ijt} on IPR_{it} would be unclear.

By contrast, the model in Section I indicates that the coefficients γ_1 and γ_2 have a clear interpretation. This is because variation in product life-cycle lengths T_j determines multinationals' sensitivity to formal patent laws, while firms' sensitivity to general institutions and development levels is theoretically independent of T_j.[31] Cross-industry variation in T_j thus allows me to capture the effect of patent laws

[29] These are of interest because without observing prices or quantities, it is unclear how observed increases in affiliate sales correspond to actual expansion in affiliate activity. Employment and the value of physical assets, however, likely reflect true changes in affiliates' manufacturing output.

[30] Prior work has established significant correlations between patent laws and general institutions, including GDP per capita and market openness (Acemoglu, Johnson, and Robinson 2005); legal origin (Lerner 2009); and economic growth (Evenson 1990). The data used in this analysis reveal a persistent correlation across countries between GDP per capita and the IPR index (68 percent in 1982, 67 percent in 2004). For further discussion, see Qian (2007).

[31] In Section V, I also verify this empirically using standard measures of general institutional quality.

separately from the effects of general institutions and development, even for cases in which all change simultaneously. Note that this identification strategy does not require identical T_j values in each country, although it is important that the ordering across industries remains relatively stable across countries.[32] This is relevant if product life cycles were to depend on local institutional or competitive environments, for example, shortening in countries with relatively weak patent laws, strong imitation capacities, or relatively intense product market competition. Because I measure T_j with US data and apply this to explain non-US outcomes, possible endogeneity of T_j with respect to host-country patent laws is unlikely to be an empirical concern.

However, other factors may contribute to higher observed levels of high-T multinational activity in strong-patent countries. Sector characteristics correlated with T_j, for example, may determine sensitivity to host-country characteristics that tend to change as patent rights improve. To accommodate this possibility, I interact key sector-level determinants of multinational activity (e.g., industrial concentration, capital intensity, R&D intensity, plant-level returns to scale) with IPR_{it} and include these as controls in a variant of (7) (Section V); note that these factors are unlikely to be nonmonotonically related to multinational activity. Relatedly, lobbying mechanisms may link multinationals' presence with the quality of local patent rights, and lobbying intensity may vary with T_j. While this form of endogeneity may be present, it is theoretically consistent with the model presented in Section I: high-T firms might have stronger incentives to encourage patent reforms, compared with firms in faster-paced industries. My empirical results would, thus, capture this additional mechanism, and their interpretation would be consistent with the model.

Finally, T_j may reflect variation across industries in the ease of imitation. For example, long product life cycles could be indicative of barriers to imitation such as product complexity; it may be that firms in short–life-cycle sectors innovate with greater intensity because their products are simpler to imitate, so that rapid innovation is a survival mechanism. Similarly, if products are well protected by patents, incumbent monopolists may have a lower (or higher) incentive to innovate relative to the case in which patents provide ineffective protection.[33] Although neither of these possibilities can be completely ruled out, the first is likely to work against finding evidence aligned with the model's results, and the second is not upheld in the data. Specifically, if the time or cost of reverse engineering were positively correlated with product lifetimes at the sector level, longer–life-cycle sectors would be less, not more, sensitive to patent laws than faster-paced sectors—the opposite of the main theoretical results. On the second point, the data show that patent effectiveness and product life-cycle lengths are not systematically related; thus, it is unlikely that long product life cycles result from barriers to imitation created by exceptionally effective patents.[34]

[32] It is straightforward to verify that this holds in an extended version of the model that allows firms to adjust T_j, at a cost, in response to the quality of local patent protection. As long as costs are not too large, firms in all industries increase T_j by the same amount, thus preserving the rank ordering across sectors.

[33] See, for example, Aghion and Howitt (1992) for discussion.

[34] The correlation between product life-cycle lengths and a standard measure of patent effectiveness from Cohen, Nelson, and Walsh (2000) is approximately 2 percent.

TABLE 3—HOST-COUNTRY PATENT LAWS AND AFFILIATE PRESENCE, INDUSTRY LEVEL

	Indicator for positive sales						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	0.0156 0.0498						
IPR × *T*		0.0611 0.0179***	1.9484 0.2766***	0.0585 0.0191***	0.9086 0.2879***	0.0994 0.0189***	0.6432 0.2609**
IPR × T^2			−0.1019 0.0152***		−0.0459 0.0161***		−0.0298 0.0146**
IPR × R&D intensity						1.6694 0.376***	1.5412 1.0024
IPR × R&D intensity^2							−1.0524 3.2375
log GDP per capita	0.1034 0.0511**						
log GDPpc × *T*				0.0017 0.0059	0.6688 0.1127***	0.0017 0.0059	0.6688 0.1127***
log GDPpc × T^2					−0.036 0.0062***		−0.036 0.0062***
Country FE, year FE, tax rate	Yes	No	No	No	No	No	No
Country-year FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,629	13,629	13,629	13,629	13,629	13,629	13,629
R^2	0.535	0.5532	0.5563	0.5532	0.5594	0.5555	0.5608

Notes: This table reports least-squares estimates of equation (7). The dependent variable indicates positive sales by affiliates of US-based multinational firms by country, sector, and year and is based on firm-level data from the BEA. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). *T* is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D Intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate. The results are robust to clustering at the sector level, excluding the top five recipients of US outward FDI, China, and India, and the chemical and pharmaceutical industries, as well as including sector-by-year fixed effects. The results shown above were estimated with OLS (Angrist and Pischke 2009), and nearly identical results obtain with probit estimation.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

IV. Main Results

As a preliminary step, I first evaluate the influence of patent laws on the broad spatial and sectoral pattern of multinational activity by estimating equation (7) with a limited dependent variable equal to one if positive affiliate sales are observed in country i and sector j during period t.

Corresponding estimates appear in Table 3. The results provide strong support for the theoretical results described above. In column 2, I find evidence that sectors with long product life cycles T_j are systematically more responsive to the strength of host-country patent protection. In column 3, I add an interaction between IPR_{it} and T_j^2, and find that the influence of patent laws follows a nonmonotonic function of T_j, reaching its highest effect in sectors with mid-length product life cycles. Columns 4 and 5 replicate columns 2 and 3 but include additional interactions between GDP per capita and T_j to better isolate the effects of patent protection from the influence of overall economic development. I find smaller but highly significant estimates under this relatively conservative approach. The estimates in column 5 imply a peak effect in an industry with $T_j = 9.90$ years (seventy-fifth percentile T_j). In columns 6 and 7, I show that the results are robust to controlling for R&D intensity.[35] Industries with high levels of R&D intensity may be relatively innovative, and thus relatively reliant on patent protection. Estimates in columns 6 and 7 reveal that R&D-intensive sectors are indeed more likely to locate affiliates in countries with strong patent laws, but controlling for this has little impact on the main coefficients of interest γ_1 and γ_2.[36] I explore the economic significance of R&D intensity relative to T_j below.

For general comparison, I also estimate the main effect of patent laws IPR_{it} on the distribution of affiliate activity, omitting interactions with product life-cycle lengths. This specification replaces η_{it} with separate country and year fixed effects η_i and η_t. The reported coefficient on the patent index in column 1 is indistinguishable from zero, revealing the potential limitations of an identification strategy that relies only on within-country time-series variation in patent laws. Finally, although all regressions shown in Table 3 are estimated using OLS, the results are nearly identical under probit estimation.

A. *Number of Multinational Affiliates, Industry Level*

Table 4 provides estimates corresponding to Result 4. The measure of offshored varieties corresponds to the measure of offshoring firms in the model; thus, I evaluate Result 4 by defining MNC_{ijt} in equation (7) to be the number of US firms with a foreign affiliate in country i and industry j during year t. From the perspective of a Southern country, this dimension of response is of interest because it is tied to potential welfare effects of patent reforms: newly introduced product manufacturing may offset the decline in local imitation while also increasing technology transfer to domestically owned firms.

Estimates appearing in Table 4 reveal strong patterns in the data that are consistent with the model in Section I. Whether or not I include interactions between GDP per capita and T_j, the estimates indicate that, conditional on the presence of multinational activity in country i, sector j, and year t, countries with stronger patent laws attract more affiliates in high-T_j sectors, with the largest effects in intermediate-T_j sectors (columns 3–7). Coefficients on the main interactions also indicate that effects of product life-cycle lengths are economically significant. Estimates in column 5 suggest that a one–standard-deviation patent reform generates (on average) a 12 percentage point differential increase in the number of affiliates between the seventy-fifth and tenth percentile industries. In addition, column 7 suggests product life-cycle lengths are a more significant determinant of sensitivity to patent laws than R&D intensity: the seventy-fifth percentile T_j industry expands 17 percentage points more than the tenth-percentile industry, compared with only

[35] R&D intensity is defined to be the industry-average R&D-to-sales ratio among sample firms.

[36] In additional robustness checks, I estimate variants of (7) in which T_j is replaced by R&D intensity and do not find significant estimates of γ_1 or γ_2. This industry characteristic does not appear to have a robust, independent effect on multinational activity (Section V).

Table 4—Host-Country Patent Laws and Number of Affiliates, Industry Level

	log number of affiliates						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	0.0822 0.0348**						
IPR × *T*		−0.0427 0.0221*	0.9145 0.3237***	0.0722 0.0356**	1.0763 0.4894**	0.1022 0.0386***	1.1164 0.4762**
IPR × T^2			−0.0504 0.017***		−0.053 0.0251**		−0.0528 0.0244**
IPR × R&D intensity						0.6671 0.3233**	1.9571 0.584***
IPR × R&D intensity^2							−5.6005 2.3403**
log GDP per capita	0.9082 0.2027***						
log GDPpc × *T*				−0.2427 0.0612***	−0.6384 0.8785	−0.241 0.0617***	−0.492 0.872
log GDPpc × T^2					0.021 0.0451		0.0135 0.0448
Country FE, year FE, tax rate	Yes	No	No	No	No	No	No
Country-year FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,783	4,783	4,783	4,783	4,783	4,783	4,783
R^2	0.7110	0.7351	0.7358	0.7387	0.7392	0.7393	0.7402

Notes: This table reports least-squares estimates of equation (7) and several alternative specifications. The dependent variable is the log of the number of affiliates of US-based multinational firms by country, sector, and year and is based on firm-level data from the BEA. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). *T* is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate. The results are robust to clustering at the sector level, Tobit estimation, negative binomial estimation, excluding the top five recipients of US outward FDI, China, and India, and the chemical and pharmaceutical industries, as well as including sector-by-year fixed effects.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

2.8 percentage points between the seventy-fifth percentile and tenth percentile industries by R&D intensity. These results also hold using negative binomial estimation.

B. *Multinational Activity, Industry Level*

To evaluate the implications of Result 4 for industry-level affiliate revenues, I define MNC_{ijt} in equation (7) to be the sales revenues earned by US firms' foreign affiliates in country *i* and industry *j* during year *t*. Estimates appear in Table 5. Consistent with the model, coefficients in panel A, column 2 indicate that the effects of patent protection follow a strong, nonmonotonic function of T_j, reaching a peak effect at $T = 9.82$ years (just below the seventy-fifth percentile of the distribution). When I include separate interactions between IPR_{it} and R&D intensity in panel A, column 3, I find nearly identical results. By contrast, the independent effect of IPR_{it}

Table 5—Host-Country Patent Laws and Affiliate Activity, Industry Level

	log affiliate sales			log affiliate assets		
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
IPR	0.0745			0.0050		
	0.0624			0.0577		
IPR × *T*		1.7241	1.7766		2.177	2.2836
		0.8375**	0.8424**		0.8481**	0.8485***
IPR × T^2		−0.0875	−0.0871		−0.112	−0.1119
		0.0434**	0.0435**		0.0441**	0.0438**
IPR × R&D intensity			2.812			4.7579
			1.3059**			1.3125***
IPR × R&D intensity^2			−8.2478			−13.2556
			5.143			5.1173**
log GDP per capita	1.7056			1.7271		
	0.3800***			0.2700***		
log GDPpc × *T*		−1.6301	−1.4267		−2.1936	−1.8253
		1.609	1.5904		1.4531	1.4228
log GDPpc × T^2		0.071	0.0606		0.0996	0.0808
		0.0828	0.0818		0.0746	0.0732
Country FE, year FE, tax rate	Yes	No	No	Yes	No	No
Country-year FE	No	Yes	Yes	No	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,783	4,783	4,783	4,783	4,783	4,783
R^2	0.6573	0.6874	0.6880	0.6600	0.6883	0.6903

(*Continued*)

is imprecisely estimated in panel A, column 1. I repeat these tests for affiliate assets in panel A, columns 4–6 and employment in panel B, columns 1–3, and find similar results.

The effects in Table 5 are also economically significant. Based on column 2 estimates, sectors with intermediate-length product life cycles (seventy-fifth percentile T_j) expand by 11 percentage points more, on average, than sectors with rapid product obsolescence (tenth-percentile T_j) following a one–standard deviation improvement in Southern patent protection. When I include interactions with R&D intensity, I find product life-cycle lengths to be more economically significant than R&D intensity by at least a factor of two. Specifically, estimates in column 3 imply that the seventy-fifth percentile T_j industry expands 21 percentage points more than the tenth-percentile industry, compared with only 8.5 percentage points between the seventy-fifth percentile and tenth-percentile industries by R&D intensity. Similar results obtain for assets and employment.

C. *Flexible Estimation*

To further evaluate the influence of patent protection on multinational activity, I estimate specifications that allow coefficients to vary flexibly across the product life-cycle length distribution. Table 6 reports estimates based on equation (8) in

TABLE 5—HOST-COUNTRY PATENT LAWS AND AFFILIATE ACTIVITY, INDUSTRY LEVEL (*Continued*)

	log affiliate employment		
Dependent variable	(1)	(2)	(3)
Panel B			
IPR	0.0178		
	0.0716		
$IPR \times T$		1.4815	1.4501
		0.7341**	0.7261**
$IPR \times T^2$		−0.0737	−0.0704
		0.0386*	0.0383*
IPR × R&D intensity			2.6333
			1.2183**
IPR × R&D intensity^2			−10.9156
			5.2993**
log GDP per capita	0.7880		
	0.472*		
log GDPpc × T		−3.1263	−3.0456
		1.6623*	1.644*
log GDPpc × T^2		0.1481	0.1441
		0.0859*	0.0851*
Country FE, year FE, tax rate	Yes	No	No
Country-year FE	No	Yes	Yes
Industry FE	Yes	Yes	Yes
Observations	4,783	4,783	4,783
R^2	0.5765	0.6141	0.6146

Notes: This table reports least-squares estimates of equation (7) and several alternative specifications. The dependent variable is the log of affiliate sales (panel A, columns 1–3), the log of affiliate assets (panel A, columns 4–6), or the log of affiliate employment (panel B, columns 1–3) for US-based multinational firms by country, sector, and year and is based on firm-level data from the BEA. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). *T* is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate. The results are robust to clustering at the sector level, Tobit estimation, excluding the top five recipients of US outward FDI, China, and India, and the chemical and pharmaceutical industries, as well as including sector-by-year fixed effects.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

which T_j is categorized by sextile (Q_1^T, Q_2^T, Q_3^T, Q_4^T, Q_5^T, and Q_6^T). A dummy corresponding to each of the top five sextiles is interacted with IPR_{it}, as follows:

$$\ln(MNC_{ijt}) = \beta + \sum_{k=2}^{6} \gamma_k \cdot IPR_{it} \times 1_{T_j \in Q_k^T} + \eta_{it} + \eta_j + \epsilon_{ijt}, \tag{8}$$

so that β and fixed effects η_{it} capture the effect of *IPR* on short–life-cycle sectors in the bottom sixth of the T_j distribution. An advantage of this approach is that the differential effect of patent laws, as reflected by the coefficients γ_2, γ_3, γ_4, γ_5, and γ_6, is unrestricted across sixths of the T_j distribution. A consistent pattern emerges from the estimates in Table 6: the third and fourth sextile coefficients tend to be positive and larger than the second, fifth, or sixth sextile coefficients, and the third sextile coefficients are significant across all specifications. This pattern indicates that expansion in multinational activity is concentrated within

Table 6—Host-Country Patent Laws and Multinational Activity, Industry Level, Flexible Estimation

	1{Positive affiliate sales}		log number of affiliates		log affiliate sales	
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
$IPR \times T_2$	−0.0019	0.0238	−0.0126	−0.001	−0.0369	−0.0242
	0.0095	0.0096**	0.047	0.0505	0.0787	0.0803
$IPR \times T_3$	0.0169	0.0485	0.1071	0.1272	0.1514	0.1759
	0.0091*	0.0093***	0.0519**	0.0584**	0.0823*	0.086**
$IPR \times T_4$	0.0607	0.0948	0.1014	0.1355	0.0852	0.1314
	0.0113***	0.0118***	0.0506**	0.0577**	0.1019	0.106
$IPR \times T_5$	0.0281	0.0563	0.0582	0.0819	−0.0041	0.0261
	0.0112**	0.0113***	0.0532	0.0577	0.0963	0.0982
$IPR \times T_6$	−0.0189	0.0189	0.0535	0.1067	0.0262	0.1025
	0.0117	0.0106*	0.0626	0.0665	0.105	0.1114
IPR × R&D intensity interactions	No	Yes	No	Yes	No	Yes
log GDPpc × *T* interactions	Yes	Yes	Yes	Yes	Yes	Yes
Country-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,629	13,629	4,783	4,783	4,783	4,783
R^2	0.5619	0.5643	0.7397	0.7404	0.6881	0.6885

	log affiliate assets		log affiliate employment	
Panel B				
$IPR \times T_2$	0.0164	0.0475	0.0228	0.0073
	0.0848	0.0871	0.078	0.0821
$IPR \times T_3$	0.1973	0.2521	0.2218	0.2085
	0.082**	0.0866***	0.0847**	0.0904**
$IPR \times T_4$	0.1169	0.2113	0.1765	0.1817
	0.0878	0.0971**	0.0943*	0.1006*
$IPR \times T_5$	−0.0468	0.0182	0.0457	0.0378
	0.0935	0.0993	0.0941	0.1038
$IPR \times T_6$	0.027	0.1754	0.1049	0.1394
	0.106	0.1166	0.1056	0.1142
IPR × R&D intensity interactions	No	Yes	No	Yes
log GDPpc × *T* interactions	Yes	Yes	Yes	Yes
Country-year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Observations	4,783	4,783	4,783	4,783
R^2	0.6895	0.6912	0.6154	0.6158

Notes: This table reports least-squares estimates of a flexible specification (8) similar to those appearing in Tables 3–5, but in which the effect of intellectual property rights is unrestricted across sextiles of the product life-cycle length distribution. The dependent variable is based on the presence (panel A, columns 1 and 2), number of affiliates (panel A, columns 3 and 4), sales (panel A, columns 5 and 6), assets (panel B, columns 1 and 2), and employment (panel B, columns 3 and 4) of US-based multinational firms based on firm-level data from the BEA. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). T_2, T_3, T_4, T_5, and T_6 indicate the second, third, fourth, fifth, and sixth sextiles of the product life-cycle length distribution, by industry, based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

intermediate-T_j industries, providing additional evidence consistent with the non-monotonicity implications of the theory.

D. *Affiliate Size and Firm Heterogeneity*

To evaluate Result 5, I investigate the effects of patent institutions on the revenues earned by individual affiliates across firms with different productivity levels by estimating an affiliate-level variant of (7)

$$\ln\left(MNC_{ijpt}\right) = \beta + \gamma_1 \cdot IPR_{it} \times T_j + \gamma_2 \cdot IPR_{it} \times T_j^2 + \eta_j + \eta_{it} + \eta_p + X_{pt} + \epsilon_{ijpt}, \tag{9}$$

where MNC_{ijpt} represents the sales of firm p's country-i, sector-j affiliates during year t. The product life-cycle length T_j corresponds to the industry reported for the affiliate. I include firm fixed effects η_p (or affiliate fixed effects η_{ijp}) as well as parent sales and R&D expenditures X_{pt} in (9) to control for differences across overall firm operations. I estimate (9) separately for affiliates within high-productivity and low-productivity firms, and also estimate a variant of (9) that includes interactions between all independent variables and an indicator for whether an affiliate is part of a low-productivity firm.[37]

The estimates in Table 7 are consistent with Result 5. Columns 1–3 and 5–7 include affiliate fixed effects, while columns 4 and 8 include firm fixed effects. In each specification, it is apparent that patent laws have a significant differential influence on affiliate sales that is concentrated within less-productive firms. Affiliates within high-productivity firms appear unaffected by patent laws relative to affiliates in low-productivity firms, which is in line with the theoretical results of Section IG that these firms are efficient enough to ensure little competition by imitators even where patent laws are weak.

E. *Exports versus Multinational Activity*

In the model, firms produce in the North and export to the South until products are relatively mature. Once a variety reaches the time-to-obsolescence cutoff $\tau^*(\xi_S)$, Southern production begins. An improvement in Southern patent protection increases this cutoff and tilts the balance of cross-border activity away from exports and toward sales by a local affiliate, with differential effects following a nonmonotonic function of T_j across industries. An alternative test of Result 4 is, thus, based on the fraction of North-to-South sales accounted for by affiliate sales.

In Table 8, I provide estimates of (7) using $MNC_{ijt} = \frac{S_{ijt}}{S_{ijt} + E_{ijt}}$, the fraction of US sales to country i in sector j during year t that are accounted for by affiliate activity, where S_{ijt} is affiliate sales and E_{ijt} is US exports. The results echo the pattern of previous estimates—a small, but significant, positive linear interaction in columns 2, 4, and 6, and a nonmonotonic curve reaching its peak at an intermediate T_j in columns 3, 5, and 7. The estimates in column 7 reveal, as before, that cross-industry

[37] To evaluate firm productivity, I compute firm-level Solow residuals and assign firms with below-median residuals to the low-productivity group. This approach relies on firm-level input and output measures, which are constructed by aggregating value added, the number of employees, and the value of physical assets (net of depreciation) across parents and affiliates within each firm. Additional details appear in the online Appendix.

Table 7—Host-Country Patent Laws and Affiliate Activity, Firm Level, Firm Heterogeneity

	log affiliate sales							
	Firm productivity				Firm productivity			
	Low	High	All firms	All firms	Low	High	All firms	All firms
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$IPR \times T$	2.1717	0.6599	−0.1039	−2.1345	2.2323	0.6055	−0.2117	−2.2251
	1.0146**	1.7042	1.6359	1.5933	1.0239**	1.7261	1.6773	1.5990
$IPR \times T^2$	−0.1145	−0.0441	−0.0042	0.1032	−0.1161	−0.0412	0.0022	0.1105
	0.0541**	0.0908	0.0873	0.0850	0.0548**	0.0921	0.0897	0.0853
$IPR \times T$ $\times$ low productivity			2.5614	3.1882			2.6982	3.2438
			1.5005*	1.4639**			1.5132*	1.4525**
$IPR \times T^2$ $\times$ low productivity			−0.1255	−0.1613			−0.1325	−0.1651
			0.0800	0.0786**			0.0810	0.0783**
$IPR \times$ R&D intensity					2.886	0.6991	1.7696	3.8733
					2.0774	1.6802	1.3945	1.1451***
$IPR \times$ R&D intensity^2					−12.3151	−3.9951	−9.0659	−18.9777
					8.3594	10.9891	9.5796	6.7845***
log GDPpc $\times$ T	−6.4591	−6.8256	−5.4754	−4.9879	−7.2241	−6.9382	−5.6412	−5.0874
	3.9507	5.7972	5.3659	2.7187*	3.949*	5.7526	5.3748	2.7259*
log GDPpc $\times$ T^2	0.3164	0.3724	0.2994	0.2622	0.3526	0.378	0.3071	0.2671
	0.2092	0.3036	0.2811	0.1453*	0.2090*	0.3017	0.2821	0.1457*
Country, year, and industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
US parent sales, US parent R&D	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Affiliate FE	Yes	Yes	Yes	No	Yes	Yes	Yes	No
Firm FE	No	No	No	Yes	No	No	No	Yes
Observations	10,517	11,988	22,505	22,505	10,517	11,988	22,505	22,505
R^2	0.7817	0.7985	0.7904	0.2889	0.7819	0.7985	0.7906	0.2901

Notes: This table reports least-squares estimates of (9), an affiliate-level analog of the main specification. Affiliates are categorized as subsidiaries of low- and high-productivity firms using a Solow residual criterion: low-productivity firms are those with a global Solow residual falling in the lower half of the distribution across firms, while all others are high-productivity firms. Each firm's Solow residual is the component of value added not explained by labor, physical assets (the value of plant, property, and equipment net of depreciation), year, and industry dummies. The dependent variable is the log of individual affiliate sales. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). *T* is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). Specifications 3, 4, 7, and 8 include all possible double and triple interactions between baseline variables and low productivity. The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

variation in T_j is approximately four times as economically significant as variation in R&D intensity, suggesting differential responses of 2.5 percentage points (T_j) and 0.8 percentage points (R&D intensity), respectively. By contrast, column 1 shows that patent laws alone are not a significant determinant of multinational activity's importance relative to exports.[38]

[38] Note that although similar measures of affiliate sales relative to exports are standard (e.g., Antras 2003; Nunn and Trefler 2008), they are only proxies. For example, US export data contain outbound intrafirm sales that are not reflected by the numerator S_{ijt}. It is also possible that affiliates produce intermediate goods that are ultimately incorporated into final goods in another sector; thus, these measures may be noisy.

TABLE 8—EXPORTS VERSUS MULTINATIONAL ACTIVITY

	Fraction of sales by multinational affiliates						
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	−0.0002 0.0092						
IPR × *T*		0.0119 0.005**	0.4533 0.0671***	0.0142 0.0054**	0.2741 0.0843***	0.0242 0.0057***	0.2021 0.0841**
IPR × T^2			−0.0238 0.0037***		−0.014 0.0047***		−0.0096 0.0047**
IPR × R&D intensity						0.386 0.0646***	0.5841 0.2259**
IPR × R&D intensity^2							−1.2363 0.8639
log GDP per capita	0.0281 0.0194						
log GDPpc × *T*				−0.0042 0.0066	0.3273 0.116***	−0.0039 0.0066	0.3277 0.1161***
log GDPpc × T^2					−0.0179 0.0064***		−0.0179 0.0064***
Country FE, year FE, tax rate	Yes	No	No	No	No	No	No
Country-year FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,651	12,651	12,651	12,651	12,651	12,651	12,651
R^2	0.5239	0.5389	0.5417	0.539	0.5425	0.5412	0.5441

Notes: This table reports least-squares estimates of equation (7) and several alternative specifications. The dependent variable is the ratio of affiliate sales to the sum of affiliate sales plus US exports by country, sector, and year for affiliates of US-based multinational firms based on firm-level data from the BEA and export data from the Census Bureau. *IPR* is the index of patent protection from Ginarte and Park (1997) and Park (2008). *T* is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate. The results are robust to clustering at the sector level, Tobit estimation, excluding the top five recipients of US outward FDI, China, and India, and the chemical and pharmaceutical industries, as well as including sector-by-year fixed effects.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

F. *The Independent Effect of Product Life-Cycle Lengths*

Result 1 indicates that patterns of offshoring by multinational firms vary across industries j depending on product life-cycle lengths T_j. To evaluate this result, I estimate the following equation:

$$MNC_{ijt} = \delta T_j + \alpha' X_j + \beta' X_{ijt} + \eta_{it} + \epsilon_{ijt},$$

where $MNC_{ijt} = \frac{S_{ijt}}{S_{ijt} + E_{ijt}}$ is the share of US affiliate sales S_{ijt} in total US sales (affiliate sales S_{ijt} plus exports E_{ijt}) to country i in sector j during year t as in Section IVE above. The coefficient δ captures the effect of product life-cycle lengths on MNC_{ijt} controlling for country-year fixed effects η_{it}, sector characteristics X_j, and

Table 9—Product Life-Cycle Lengths and Offshoring Intensity

	Fraction of sales by multinational affiliates				
Dependent variable	(1)	(2)	(3)	(4)	(5)
T	−0.0243	−0.1474	−0.0984	−0.2495	−0.1813
	0.0064***	0.0509***	0.0506*	0.0522***	0.0509***
log GDPpc × T		0.0137	0.0115	0.0283	0.0271
		0.006**	0.006*	0.0061***	0.0061***
R&D intensity			1.0883	−4.1162	−3.5703
			0.1102***	0.7229***	0.7166***
log GDPpc × R&D intensity				0.5842	0.5227
				0.0836***	0.0821***
Country-year FE	Yes	Yes	Yes	Yes	Yes
Industry characteristics	No	No	No	No	Yes
Observations	12,651	12,651	12,651	12,651	11,928
R^2	0.3285	0.3290	0.3436	0.3482	0.4668

Notes: This table reports least-squares estimates corresponding to a first-order test of the prediction that the share of offshored economic activity by industry declines in product life-cycle lengths. The dependent variable is the ratio of affiliate sales to the sum of affiliate sales plus US exports by country, sector, and year, for affiliates of US-based multinational firms based on firm-level data from the BEA and export data from the Census Bureau. T is the product life-cycle length, by industry, and is the average patent citation lag based on data from the USPTO and NBER. R&D intensity is the average ratio of R&D to sales by industry based on BEA data, and GDP per capita (GDPpc) is from the Penn World Table, Heston, Summers, and Aten (2009). Industry characteristics in column 5 are patent and secrecy effectiveness (Cohen, Nelson, and Walsh 2000), capital and labor intensity, plant-level returns to scale, and the Herfindahl index; note that coverage of patent and secrecy effectiveness is incomplete, limiting the sample size in column 5 relative to columns 1–4. The sample period is 1982–2004. Standard errors, adjusted for clustering at the country level, appear below each point estimate.
***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

country-sector-year covariates X_{ijt}.[39] Estimation results appear in Table 9. Consistent with Result 1, I find that the estimated coefficient δ is negative and highly significant, ranging from between -0.024 (standard error $= 0.0064$) to -0.18 (standard error $= 0.051$) across specifications. These effects are economically significant, indicating that a one–standard deviation difference across industries in T_j explains a 2–14 percentage point difference in the fraction of US sales accounted for by affiliate sales.

V. Robustness Checks and Alternative Specifications

I subject the main results to a number of robustness checks to better establish their stability. First, I reevaluate all main results with a full set of sector-year fixed effects; these control for the possibility that entire industries may undergo significant changes in size, competitiveness, or innovativeness over time. I estimate specifications that control for the possibility that other sector characteristics may influence

[39] Industry characteristics X_j include R&D intensity, capital intensity, labor intensity, plant-level fixed costs, industrial concentration, patent effectiveness, and secrecy effectiveness. Country-sector-year covariates X_{ijt} include interactions between GDP per capita and T_j and between GDP per capita and R&D intensity.

sensitivity to patent laws by including separate interactions between IPR_{it} and sector-level measures of industrial concentration, labor intensity, capital intensity, plant fixed costs, and patent and secrecy effectiveness (online Appendix Table A.1). In separate tests, I exclude the chemical and pharmaceutical industries to ensure that the findings are general and not specific to industries for which patents are known to be exceptionally effective (Levin et al. 1987; Cohen, Nelson, and Walsh 2000) and exclude the top five recipients of US outward foreign direct investment. To reduce the influence of measurement error in product life-cycle lengths, I cluster standard errors by sector. In addition, I reestimate Tables 3–8 using a Tobit strategy that allows for left-censoring at zero. For each of these tests, I find qualitatively similar results.[40]

I perform additional tests using alternative product life-cycle length indexes, described in Section II. First, I find nearly identical results using the seventy-fifth and eighty-fifth percentile of patent citations instead of my main measure of T (the average). This ensures that my sector measure is not driven by unsuccessful innovations that receive few citations. Tests based on patent-renewal data also indicate that long–life-cycle sectors are significantly more sensitive to patent laws. By contrast, separate results suggest that the Broda and Weinstein (2010) product turnover index is not a strong predictor of firms' sensitivity to patent laws. This latter finding is in line with the observation that retail-level product turnover reflects characteristics beyond the broader product life cycle with which firms facing imitation risk associated with their manufacturing location are concerned.[41]

A. *First Differences*

While Result 4 concerns the pattern of multinational activity across countries with varied patent regimes, Results 2 and 3 provide qualitatively similar statements regarding how affiliate activity responds to patent reform within a country. To evaluate these results, I estimate a first-differenced version of (7)

$$(10) \quad \Delta MNC_{ijt} = \gamma_1 \cdot \Delta IPR_{it} \times T_j + \gamma_2 \cdot \Delta IPR_{it} \times T_j^2 + \Delta \eta_{it} + \Delta \epsilon_{ijt},$$

where ΔMNC_{ijt} is an indicator for increased affiliate sales. Note that the constant term and sector fixed effect have dropped out of the regression equation, while the country-year fixed effect is now replaced by a differenced $\Delta\eta_{it}$. Corresponding results appear in Table A.2 and provide confirmation for Result 3. Estimates in columns 1 though 7 reveal a qualitatively similar pattern of sensitivity to reforms compared with columns 1–7 in Table 3; the estimates are also of comparable magnitude

[40] In addition, I use the within-industry variation in $\hat{T}_j$ to assess whether my main proxy is more informative for low-variance sectors and find that it is: the results hold with greater strength among three-digit industries with below-median variation in $\hat{T}_j$ across patent classes. I also perform tests based on the subcomponents of the patent protection index and find that membership in international intellectual property treaties (which impose external enforcement mechanisms) is the most important determinant of multinational activity. In addition, IPR_{it} likely captures de facto patent enforcement imperfectly; I therefore perform a series of additional tests in which I replace IPR_{it} with (i) the treaty-membership subindex, or (ii) an interaction between IPR_{it} and standard measures of institutional quality: *Rule of law* and *Corrupt*, both from La Porta et al. (1998). The empirical results described in Section IV are robust to either approach.

[41] One important caveat to results based on the Broda and Weinstein (2010) product turnover index involves industrial composition: the sectoral overlap between this index and multinational activity is limited (20–25 percent), so that these results may not be compared directly with my main estimates.

and significance. The results are nearly identical when ΔMNC_{ijt} is defined to be an indicator for increased affiliate assets or employment. I also estimate the same set of specifications with ΔMNC_{ijt} defined to be the one-period difference in affiliate sales. In this latter set of results, I find that the signs match the theory in each case, although the key interaction coefficients are significant only in corresponding columns 2, 3, and 6.

B. *Patent Effectiveness*

Industry surveys (Levin et al. 1987; Cohen, Nelson, and Walsh 2000) have revealed substantial cross-industry variation in the effectiveness of patents and secrets as means of protecting innovations. This is important, because the model presumes that host-country patent laws offer effective protection from imitation where enforced. The surveys mentioned above suggest, however, that innovations in certain sectors in fact receive little protection, even in developed countries with strong patent laws. I incorporate this observation into my analysis using measures of product patent- and secrecy effectiveness from Cohen, Nelson, and Walsh (2000).

In Table A.3, I present estimates categorized by whether a sector is above or below the median score for product patent effectiveness, where high scores indicate high effectiveness. As expected, the mechanism of patent law sensitivity predicted by my model operates more forcefully in sectors for which patents are effective. I show results for each of two dependent variables (log affiliate sales and log number of affiliates). Columns 3 and 6 are triple-differenced specifications that include interactions between IPR_{it}, T_j, or T_j^2, and patent effectiveness; the coefficient on the linear (quadratic) interaction is positive (negative) and highly significant in these specifications. In unreported results, I take a similar approach using secrecy effectiveness and reach a similar conclusion: sectors for which secrecy is ineffective tend to be more sensitive to host-country patent laws.

VI. Conclusion

This article has examined multinationals' sensitivity to host-country intellectual property institutions both theoretically and empirically. Within a model of firms' global production location decisions, I develop results regarding the spatial and sectoral composition of multinational activity. The model suggests that sensitivity to local patent institutions is concentrated among sectors with relatively long product life cycles, with the most pronounced sensitivity in sectors with mid-length product life cycles. Among sensitive sectors, stronger host-country patent laws attract a relatively larger number of affiliates and lead existing affiliates to expand.

These results find robust empirical support within a comprehensive panel of US multinationals' activity spanning 92 countries and 37 industries during 1982–2004. Using the interaction between patent laws and product life-cycle lengths, I am able to explain systematic variation in affiliate activity, as measured by sales, assets, and employment. The results provide evidence that cross-industry differences in the rate of product obsolescence are a significant determinant of firms' sensitivity to host-country patent laws and establish the causal effect of patent laws on multinational activity.

My findings speak to an ongoing debate over the extent to which developing countries should protect intellectual property. Strengthened patent protection may discourage imitation but raises prices faced by domestic consumers, creating direct welfare losses such as those found by Chaudhuri, Goldberg, and Jia (2006). Unless these losses are eventually offset by higher growth, for example due to significant increases in domestic innovation and technology transfer by foreign firms, it is not clear that developing countries stand to gain by undertaking the costly investment of improving patent protection. My results reveal that stronger patents do attract multinational activity, itself an important channel for technology transfer, but primarily in the subset of sectors with relatively long-lived intellectual property. In addition, if stronger patent rights also attract greater levels of arms'-length licensing, my estimates may understate the overall effect of local patent protection on technology transfer.

Finally, the theoretical and empirical results of this article are consistent with the idea that technology obsolescence rates condition firms' reliance on formal intellectual property protection. Moreover, while previous work has identified strategic substitutes for patent protection such as secrecy, lead time, and product complexity (Moser 2005; Cohen et al. 2000), this article provides evidence—although indirect—that rapid product innovation may itself form an additional strategic substitute for patent protection. From a policy perspective, this latter observation raises the further possibility that formal patent protection may therefore be especially important for products embodying long-lived technologies. Further investigation of these possibilities and their implications for patent policy is an important area for future research.

REFERENCES

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2005. "Institutions as a Fundamental Cause of Long-Run Growth." In *Handbook of Economic Growth*. Vol. 1A, edited by Philippe Aghion and Steven N. Durlauf, 385–472. Amsterdam: Elsevier B.V.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–51.

Antràs, Pol. 2003. "Firms, Contracts, and Trade Structure." *Quarterly Journal of Economics* 118 (4): 1375–1418.

Antràs, Pol. 2005. "Incomplete Contracts and the Product Cycle." *American Economic Review* 95 (4): 1054–73.

Antràs, Pol, Mihir A. Desai, and C. Fritz Foley. 2009. "Multinational Firms, FDI Flows, and Imperfect Capital Markets." *Quarterly Journal of Economics* 124 (3): 1171–1219.

Antràs, Pol, and Elhanan Helpman. 2004. "Global Sourcing." *Journal of Political Economy* 112 (3): 552–80.

Arkolakis, Costas, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple. 2013. "Innovation and Production in the Global Economy." National Bureau of Economic Research Working Paper 18972.

Arora, Ashish. 1996. "Contracting for Tacit Knowledge: The Provision of Technical Services in Technology Licensing Contracts." *Journal of Development Economics* 50 (2): 233–56.

Arora, Ashish, Marco Ceccagnoli, and Wesley M. Cohen. 2003. "R&D and the Patent Premium." National Bureau of Economic Research Working Paper 9431.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. 2010. "Intrafirm Trade and Product Contractibility." *American Economic Review* 100 (2): 444–48.

Bessen, James. 2008. "The Value of U.S. Patents by Owner and Patent Characteristics." *Research Policy* 37 (5): 932–45.

Bilir, L. Kamran. 2014. "Patent Laws, Product Life-Cycle Lengths, and Multinational Activity: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.104.7.1979.

Bils, Mark. 2009. "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *Quarterly Journal of Economics* 124 (2): 637–75.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does Management Matter? Evidence from India." *Quarterly Journal of Economics* 128 (1): 1–51.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.

Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and Explaining Management Practices across Firms and Countries." *Quarterly Journal of Economics* 122 (4): 1351–1408.

Branstetter, Lee G., Raymond Fisman, and C. Fritz Foley. 2006. "Do Stronger Intellectual Property Rights Increase International Technology Transfer? Empirical Evidence from U.S. Firm-Level Panel Data." *Quarterly Journal of Economics* 121 (1): 321–49.

Branstetter, Lee, Raymond Fisman, C. Fritz Foley, and Kamal Saggi. 2011. "Does Intellectual Property Rights Reform Spur Industrial Development?" *Journal of International Economics* 83 (1): 27–36.

Broda, Christian, and David E. Weinstein. 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review* 100 (3): 691–723.

Chaudhuri, Shubham, Pinelopi K. Goldberg, and Panle Jia. 2006. "Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India." *American Economic Review* 96 (5): 1477–1514.

Cohen, Wesley M., Richard R. Nelson, and John P. Walsh. 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (Or Not)." National Bureau of Economic Research Working Paper 7552.

Criscuolo, Chiara, Jonathan E. Haskel, and Matthew J. Slaughter. 2010. "Global Engagement and the Innovation Activities of Firms." *International Journal of Industrial Organization* 28 (2): 191–202.

Davidson, W. H., and Donald G. McFetridge. 1985. "Key Characteristics in the Choice of International Technology Transfer Mode." *Journal of International Business Studies* 16: 5–21.

Doms, Mark E., and J. Bradford Jensen. 1998. "Comparing Wages, Skills, and Productivity between Domestically and Foreign-Owned Manufacturing Establishments in the United States." In *Geography and Ownership as Bases for Economic Accounting.* National Bureau of Economic Research Studies in Income and Wealth. Vol. 59, edited by Robert E. Baldwin, Robert E. Lipsey, and J. David Richards, 235–58. Chicago: University of Chicago Press.

Evenson, Robert E. 1990. "Survey of Empirical Studies." In *Strengthening Protection of Intellectual Property in Developing Countries: A Survey of the Literature*, edited by Wolfgang E. Siebeck, 33–46. World Bank Discussion Papers. Washington, DC: World Bank.

Ginarte, Juan C., and Walter G. Park. 1997. "Determinants of Patent Rights: A Cross-National Study." *Research Policy* 26 (3): 283–301.

Glass, Amy Jocelyn, and Kamal Saggi. 2002. "Licensing versus Direct Investment: Implications for Economic Growth." *Journal of International Economics* 56 (1): 131–53.

Greenstein, Shane M., James B. Wade. 1998. "The Product Life Cycle in the Mainframe Computer Market, 1968–1982." *RAND Journal of Economics* 29 (4): 772–89.

Grossman, Gene M., and Elhanan Helpman. 1991. "Quality Ladders and Product Cycles." *Quarterly Journal of Economics* 106 (2): 557–86.

Grossman, Gene M., and Edwin L. C. Lai. 2004. "International Protection of Intellectual Property." *American Economic Review* 94 (5): 1635–53.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." National Bureau of Economic Research Working Paper 8498.

Helpman, Elhanan. 1993. "Innovation, Imitation, and Intellectual Property Rights." *Econometrica* 61 (6): 1247–80.

Heston, Alan, Robert Summers, and Bettina Aten. 2009. *Penn World Table Version 6.3*. Philadelphia: Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

Igami, Mitsuru. 2010. "Who Offshores, Where, and How? Evidence from the Hard Disk Drive Industry." Unpublished.

Javorcik, Beata Smarzynska. 2004. "The Composition of Foreign Direct Investment and Protection of Intellectual Property Rights: Evidence from Transition Economies." *European Economic Review* 48 (1): 39–62.

Keller, Wolfgang, and Stephen R. Yeaple. 2013. "The Gravity of Knowledge." *American Economic Review* 103 (4): 1414–44.

Klepper, Steven. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *American Economic Review* 86 (3): 562–83.

Klepper, Steven, and Sally Sleeper. 2005. "Entry by Spinoffs." *Management Science* 51 (8): 1291–1306.

Krugman, Paul. 1979. "A Model of Innovation, Technology Transfer, and the World Distribution of Income." *Journal of Political Economy* 87 (2): 253–66.

La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert W. Vishny. 1998. "Law and Finance." *Journal of Political Economy* 106 (6): 1113–1155.

Lanjouw, Jean O., Ariel Pakes, and Jonathan Putnam. 1998. "How to Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data." *Journal of Industrial Economics* 46 (4): 405–32.

Lanjouw, Jean O., and Mark Schankerman. 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators." *Economic Journal* 114 (495): 441–65.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter. 1987. "Appropriating the Returns from Industrial Research and Development." *Brookings Papers on Economic Activity* 18 (3): 783–832.

Lerner, Josh. 2009. "The Empirical Impact of Intellectual Property Rights on Innovation: Puzzles and Clues." *American Economic Review* 99 (2): 343–48.

Manova, Kalina, Shang-Jin Wei, and Zhiwei Zhang. 2011. "Firm Exports and Multinational Activity Under Credit Constraints." National Bureau of Economic Research Working Paper 16905.

Mansfield, Edwin, Mark Schwartz, and Samuel Wagner. 1981. "Imitation Costs and Patents: An Empirical Study." *Economic Journal* 91 (364): 907–18.

McCalman, Phillip. 2004. "Foreign Direct Investment and Intellectual Property Rights: Evidence from Hollywood's Global Distribution of Movies and Videos." *Journal of International Economics* 62 (1): 107–23.

Mehta, Aditi, Marc Rysman, and Tim Simcoe. 2010. "Identifying the Age Profile of Patent Citations: New Estimates of Knowledge Diffusion." *Journal of Applied Econometrics* 25 (7): 1179–1204.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.

Moser, Petra. 2005. "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs." *American Economic Review* 95 (4): 1214–36.

Narin, Francis, and Dominic Olivastro. 1993. "Patent Citation Cycles." *Library Trends* 41 (4): 700–09.

National Science Board. 2010. *Science and Engineering Indicators 2010*. Arlington: National Science Foundation (NSB 10-01).

Nunn, Nathan, and Daniel Trefler. 2008. "The Boundaries of the Multinational Firm: An Empirical Analysis." In *The Organization of Firms in a Global Economy*, edited by Elhanan Helpman, Dalia Marin, and Theirry Verdier, 55–83. Cambridge: Harvard University Press.

Pakes, Ariel S. 1986. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica* 54 (4): 755–84.

Park, Walter G. 2008. "International Patent Protection: 1960–2005." *Research Policy* 37 (4): 761–66.

Poole, Jennifer P. 2013. "Knowledge Transfers from Multinational to Domestic Firms: Evidence from Worker Mobility." *Review of Economics and Statistics* 95 (2): 393–406.

Qian, Yi. 2007. "Do National Patent Laws Stimulate Domestic Innovation in a Global Patenting Environment? A Cross-Country Analysis of Pharmaceutical Patent Protection, 1978–2002." *Review of Economics and Statistics* 89 (3): 436–53.

Qian, Yi. 2008. "Impacts of Entry by Counterfeiters." *Quarterly Journal of Economics* 123 (4): 1577–1609.

Sampat, Bhaven N. 2010. "When Do Patent Applicants Search for Prior Art?" *Journal of Law and Economics* 53 (2): 399–416.

Schankerman, Mark. 1998. "How Valuable Is Patent Protection? Estimates by Technology Field." *RAND Journal of Economics* 29 (1): 77–107.

Schankerman, Mark, and Ariel Pakes. 1986. "Estimates of the Value of Patent Rights in European Countries during the Post-1950 Period." *Economic Journal* 96 (384): 1052–76.

Teece, David J., Gary Pisano, and Amy Shuen. 1997. "Dynamic Capabilities and Strategic Management." *Strategic Management Journal* 18 (7): 509–33.

United Nations Conference on Trade and Development. 2005. *World Investment Report 2005: Transnational Corporations and the Internationalization of R&D*. New York: United Nations.

Vernon, Raymond. 1966. "International Investment and International Trade in the Product Cycle." *Quarterly Journal of Economics* 80 (2): 190–207.

Yang, Guifang, and Keith E. Maskus. 2001. "Intellectual Property Rights, Licensing, and Innovation in an Endogenous Product-Cycle Model." *Journal of International Economics* 53 (1): 169–87.