Patent Laws, Product Lifecycle Lengths, and the Global Sourcing Decisions of U.S. Multinationals

L. Kamran Bilir*

July 5, 2011

Abstract

This paper investigates the influence of patent laws on firms' global sourcing decisions. I develop a theoretical model of multinational firms' location and production decisions in the presence of cross-country differences in intellectual property rights and cross-sector differences in the length of product lifecycles. I show that patent reforms are irrelevant to firms' sourcing decisions in industries with rapid product turnover. By contrast, strong patent laws attract affiliate activity in industries with longer product lifecycles, because products in these industries are more likely to be imitated prior to obsolescence and are thus more reliant on patent enforcement to protect revenues. These effects are more pronounced for less-productive firms. Using comprehensive panel data on the sales, assets, and employment of U.S. multinationals and their affiliates abroad and a new measure of product obsolescence, I find robust empirical support for these predictions. Effects are significant along all margins of multinational activity, including multinational presence by country and sector, total affiliate sales conditional on presence, the number of affiliates, and affiliate-level sales. In addition, I find that stronger patent rights tilt the balance of cross-border activity away from exports and toward multinational activity. Finally, my identification strategy allows me to isolate the causal effect of patent reforms on multinational operations, which the prior literature has struggled to establish because of concurrent policy reforms.

*Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072, kbilir@stanford.edu. The statistical analysis of firm-level data on U.S. multinational companies was conducted at the Bureau of Economic Analysis, U.S. 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 U.S. Department of Commerce. I thank Kenneth Arrow, Nick Bloom, Kalina Manova, Petra Moser, and Bob Staiger for their invaluable guidance. I also thank Dietmar Harhoff, Mark Lemley, Mitch Polinsky, Stanley Watt, and Bill Zeile for insightful conversations, and seminar participants at Stanford University for 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.

1 Introduction

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 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 sourcing decisions in all sectors equally. Consider for example the experience of two firms in the electronics and solar cell industries, Apple and Solar Junction. Apple chooses to produce even its newest products in locations with weak intellectual property institutions. On the other hand, Solar Junction, a U.S. firm that has developed high-efficiency solar cells, is building production capacity for its latest models in the United States. Although these two firms differ along many dimensions, Solar Junction points to concerns about imitation risk as a major factor behind the decision. In particular, it cites the durability of the intellectual property associated with its current products—which is long-lived compared with Apple's—as a key underlying cause for its sensitivity to intellectual property laws.[3]

This paper provides evidence that multinationals' sensitivity to host-country intellectual property protection is determined by the length of product lifecycles. I develop a global sourcing model in which innovating firms in the North face imitation risk in the South. The quality of intellectual property protection differs across countries and, though all products eventually become obsolete, product lifecycle lengths vary across sectors. The model offers detailed predictions for the location and sectoral composition of multinational activity which find robust empirical support in a comprehensive dataset on U.S. multinational firms.

In the model, innovating firms make production and location decisions for products with industry-specific lifecycle lengths that are technologically determined. Production requires both a headquarters input and a mobile manufacturing input, and the latter may be located in either the North or the South. Firms enjoy lower wages when manufacturing in the South, but patents there are poorly 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

[1] Criscuolo, Haskel, and Slaughter (2010), Doms and Jensen (1998), UNCTAD (2005).
[2] National Science Board (2010).
[3] Based on conversations with a senior research engineer at Solar Junction.

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 lifecycle.

I show that firms follow a sector-invariant optimal sourcing rule, whereby production moves to the South when products reach a critical time-to-obsolescence cut-off. Improvements to patent protection in the South increase this cut-off. Because products with lifetimes shorter than the cut-off will always be manufactured in the South, patent reforms will have no effect on sourcing decisions in fast-turnover industries. By contrast, the manufacture of longer-lived products will be offshored to the South earlier in the lifecycle following the reform, increasing multinational activity in these sectors. Moreover, the response to patent reforms is a non-monotonic function of product lifecycle lengths, and is most pronounced in intermediate lifecycle length sectors. Intuitively, this is because the increase in the sourcing cut-off 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 enforcement. Finally, I show that all of these effects are concentrated among relatively less-productive firms.

I find strong support for these predictions in a panel of affiliate-level data on U.S. multinational firms from the Bureau of Economic Analysis. The dataset spans 92 countries and 37 industries during 1982–2004. I empirically evaluate how interactions between a country-level index of patent strength and a sector-level measure of product lifecycle lengths predict patterns of multinational activity. I explicitly control for differences across sectors in product complexity (R&D intensity), which affects imitation costs, as well as for host-country corporate tax rates. Using cross-sector variation in product lifecycle lengths, I am able to identify the influence of patent laws separately from that of other host-country institutions and overall economic development as proxied by GDP per capita.

To test the model predictions, I develop a new measure of product lifecycle lengths using information contained in patent records from the NBER U.S. Patent Citations Data File. The measure is based on the idea that the duration of citations received by patents reflects the lifetime of technologies embodied in products. I construct the average citation lag for each patent within over 400 unique U.S. patent classes, and use moments of the class-specific average lag distributions as sector-level proxies for product lifecycle lengths. Host-country patent laws are measured with a widely-used index of patent protection developed by Ginarte and Park (1997) and extended by Park (2008).

Figures 1 and 2 offer motivating evidence that product lifecycle lengths significantly in-

fluence the sensitivity of multinational activity to host-country patent laws. The figures compare the distribution of affiliate sales between countries with weak and strong patent regimes (Figure 1) and within reforming countries over time (Figure 2). It is apparent that patent reforms induce a significant compositional shift across sectors with different product lifetimes: fast-turnover sectors account for a smaller share of affiliate sales when patent laws are relatively strong, while intermediate sectors account for a larger share.

Five main results emerge from the empirical analysis. First, relative to fast-turnover sectors, sectors with long product lifetimes are significantly more likely to locate affiliate activity in countries with strong patent protection than in countries with weak patent protection. Second, conditional on hosting any multinational activity in a given sector, levels of affiliate sales, assets, and employment respond to patent strength according to the sector's product lifecycle length. In particular, sensitivity to patent protection follows a curve that is low in short-lifecycle sectors and high in long-lifecycle sectors, with the largest effects in sectors near the 75th percentile of the distribution.

Third, separating the industry-level responses along the extensive and intensive margins, I find that the pace of product turnover is again an important determinant of both. Consistent with the predictions of my model, stronger patent laws both attract more affiliates (extensive margin) and expand the size of existing affiliates (intensive margin) in sectors with long product lifecycles, with the largest effects in sectors with mid-length lifecycles. This finding is important from a Southern welfare perspective, because it reveals that stronger patents attract new affiliates, expanding not just the level but also the scope of local industrial activity. This latter effect strongly suggests that better patent laws attract greater levels of technology transfer in long-lifecycle sectors. Fourth, I verify that firm-level responsiveness to patent laws is more pronounced among the affiliates of low-productivity parents. Finally, combining data on U.S. exports and multinational activity, I show that patent reforms tilt cross-border activity away from exports and toward multinational activity in sectors with relatively long product lifetimes.

These results suggest that incorporating heterogeneity in product lifecycle lengths across industries can rationalize the systematic variation in multinational activity across sectors with different rates of product obsolescence and across countries with different levels of patent protection. This differential sensitivity to patent laws is also economically significant. A one standard deviation improvement in measured patent protection attracts between 10 and 20 percentage points more multinational activity in the 75th-percentile sector than in the 10th-percentile sector by product lifecycle lengths. Similarly, the effect at the 75th-percentile

exceeds that at the 95th-percentile sector by nearly 10 percentage points. These magnitudes are two to four times larger than comparative statics for other sector-level determinants of multinationals' sensitivity to patent reforms such as R&D intensity, and hold after accounting for cross-sector differences in labor and capital intensities as well as plant-level returns to scale.

My paper 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 a dummy for a sector's technology intensity and cross-section survey 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 (2010) 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 U.S. 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.[4] I build on the foundation of these prior analyses by introducing systematic, continuous variation across sectors in product lifecycle lengths, a dimension that determines sensitivity to local patent laws but not to general institutions or economic development. This variation enables me to isolate patent laws' casual influence on multinational activity, even in the presence of concurrent policy reforms.[5]

The theoretical model in this paper 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

[4]As will be described in section 4, measuring the effect of patent laws in a standard regression set-up will result in estimates reflecting multinationals' responsiveness not only to patent laws, but also to the quality of other institutions and levels of economic development. Importantly, binary categorizations such as patent-intensity or technology-intensity do not mitigate this concern, because high-patent or high-technology firms are likely to be more reliant on factors that are relatively abundant in countries with higher levels of general economic development, such as skilled labor.

[5]Bilir, Moser, and Talis (2010) apply a different strategy to identify the impact of the Paris Convention, the first international patent treaty, on technology transfer to the United States. Rather than using countries' individual accession dates, which are correlated with the timing of domestic reforms, they use U.S. accession and find that existing members disproportionately increased U.S. patenting in response to strengthened rights.

(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.

My results complement prior studies that have investigated differences across industries in the importance and effectiveness of patent protection. Firm-level surveys (Mansfield 1981, Levin et al 1987, and Cohen, Nelson, and Walsh 2000) and other analyses (Schankerman 1998, Harhoff 2000, Arora et al 2003, and Moser 2003) reveal large differences in the effectiveness of patents as a means of appropriating the returns from innovation, with patents conferring exceptionally effective protection in the chemical and pharmaceutical industries.[6] 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.

My work also relates to an emerging literature that documents Schumpeterian creative destruction at the product level. Analyzing entry and exit rates among household and consumer goods, Broda and Weinstein (2010) discover substantially higher rates of product turnover within knowledge-intensive sectors relative to non-technical sectors. Bils and Klenow (2004) find similar evidence based on the frequency of price changes in the Consumer Price Index across sectors due to discontinued products.[7] By contrast, I focus on product obsolescence more broadly defined, since the intellectual property relevant for imitation may span multiple versions of a product during its lifecycle.

Finally, this paper contributes to a line of research examining the impact of institutional frictions on foreign direct investment. Recent empirical 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 2010; Bernard, Jensen, Redding, and Schott 2010; Antràs, Garicano, and Rossi-Hansberg 2008; Antràs 2003). Others have suggested that the effects of these imperfections are acutely felt by

[6]Numerous factors have been proposed as contributors to the observed cross-industry differences in patent effectiveness, including the ability for competitors to legally circumvent patents, the complexity of products, the cost of product development, and the pace of technological change. In this paper, I emphasize the last of these factors, building on observations made by previous innovation scholars (Moser 2003, Schankerman and Pakes 1986) that the majority of patented products become obsolete well before associated patents expire. These observations imply that the duration of a patent-based monopoly is, on average, an industry characteristic rather than a uniform legal standard. See also Burk and Lemley (2009) for an excellent discussion and review of this evidence.

[7]In addition, cross-industry differences in product churning and trade flow duration are also documented in Manova (2007) and Besedes and Prusa (2006a, 2006b).

innovative firms, particularly those seeking to manufacture cutting-edge technology abroad (Antràs and Helpman 2004, Nunn and Trefler 2008, Antràs 2005, Davidson and McFetridge 1985).

The rest of the paper presents my theoretical and empirical analysis. In section 2, I develop a global sourcing model with innovating firms that face imitation risk. After describing the data in section 3, I outline the estimation approach that will be used to test the model's predictions in section 4. In sections 5 and 6, I describe the empirical results. Section 7 concludes.

2 Motivating Theory

The model developed below investigates the global sourcing trade-off between imitation risk and production costs from the perspective of a multinational firm. Specifically, the model formalizes the idea that this trade-off evolves over the product lifecycle, and that sourcing decisions thus hinge on both host-country characteristics and the rate of product obsolescence. To highlight the theoretical predictions that I will test in the empirical analysis to follow, I present a partial-equilibrium analysis with two countries. At the end of the section, I show how the results from this baseline model extend to settings with multiple Southern countries and heterogeneous firms.

2.1 Demand

There are J sectors, each with a continuum of differentiated product varieties. Aggregate consumption of sector-j goods is D_j, and consumers in all countries share the following preferences over varieties in the J sectors

$$U = \prod_j D_j^{\mu_j}, \quad \sum_j \mu_j = 1, \quad D_j = \left(\int_0^{M_j} d_{jk}^{\rho} dk \right)^{1/\rho}, \quad \rho \equiv \frac{\sigma - 1}{\sigma} < 1, \quad \sigma > 1, \tag{1}$$

where M_j is the measure of varieties produced in sector j. A property of this preference function is that consumers will spend a fixed fraction of expenditures on each type of good. This, combined with the consumer's optimality conditions, implies that demand for a sector-

j, variety-k good will be

$$d_{jk} = D_j \left(\frac{p_{jk}}{P_j}\right)^{-\sigma} = p_{jk}^{-\sigma} \left(\frac{\mu_j Y}{P_j^{1-\sigma}}\right) \equiv p_{jk}^{-\sigma} \lambda_j \tag{2}$$

where $P_j \equiv \left[\int_0^{M_j} p_{jk}^{1-\sigma} dk\right]^{\frac{1}{1-\sigma}}$ is the price index for industry j and Y is total expenditure. To highlight the effects of cross-sector variation in product lifecycle lengths, I assume preferences are symmetric across industries, $\mu_j = \mu$, $j = 1, ..., J$.[8]

2.2 Production

Time is continuous. At every moment, firms in each sector produce differentiated varieties k and sell to consumers in both the North and the South, symmetric countries that are each of size 1.[9] I assume innovating firms pay a one-time fixed cost f_j to develop a new variety in sector j, thereafter enjoying a monopoly until the variety is either imitated or becomes obsolete.[10] Sectors are distinguished by the pace of product obsolescence, which I assume is determined by exogenous underlying technological developments specific to each industry.[11] Product lifecycle lengths T_j thus vary across sectors, but are shared by all sector-j products. 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. Any intellectual property and imitation products associated with the retired variety also immediately become obsolete. This approach builds on observations made by previous innovation scholars (Moser 2003, Schankerman and Pakes 1986) that the majority of patented products become obsolete well before associated patents expire, implying that the

[8] As usual, in the background of this partial equilibrium setting is a traded numeraire good that enters quasi-linearly into utility and is always consumed in positive amounts. See also note 9.

[9] These symmetry assumptions are made for simplicity. Assuming instead that the Northern market is larger than the Southern market will not affect the qualitative predictions of the model, provided that consumers in both countries demand at least a small amount of each variety. However, a related possibility is that relative market sizes vary across industries (i.e. preferences are non-homothetic). While beyond the scope of the present model, I include sector (or sector-year) fixed effects in my empirical analysis to ensure that my results are robust to this consideration.

[10] Similar to Grossman and Helpman (1991), I assume that any two firms producing the same variety engage in Bertrand competition. As a result, at most one Northern firm will produce a given variety. Southern imitators, on the other hand, will enjoy lower marginal costs than innovating firms. Circumstances will thus arise in which pursuing imitation is an economically attractive investment for Southern entrepreneurs.

[11] What is important for my empirical analysis is that the product lifecycle lengths of U.S. firms do not respond significantly to changes in the patent laws of foreign countries. Taking product lifecycle lengths 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.

duration of a patent-based monopoly is, on average, an industry characteristic rather than a uniform legal standard.

To keep things simple, I also 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, M_j, which I normalize to 1 for all sectors j.[12] I will discuss the implications of endogenous entry in later notes, but since my aim is to evaluate the sourcing decisions of innovating firms in response to the strength of Southern patent laws, this assumption is reasonable provided that Northern multinationals make innovation decisions primarily on the basis of Northern patent laws. Finally, notice that under my assumptions, the distribution of product maturities will be uniform with density $\psi_j(t) \equiv 1/T_j$ in each sector j.

Innovating Firms

Consider a sector-j firm that is able to produce a particular variety k. To produce, the firm combines headquarters services x_h and manufacturing x_l according to a Cobb-Douglas production technology $q = Ax_h^{\alpha}x_l^{1-\alpha}$, where $A = \alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)}$ for convenience. Both inputs require a unit of labor, but only the manufacturing input is mobile—innovating firms locate permanent headquarters in the North and thus source x_h in the North,[13] but at any time may choose to costlessly shift manufacturing activity x_l to an affiliate in the South, where wages are lower: $w^S < w^N$ where w^S denotes Southern wages and w^N denotes Northern wages.[14] Final goods are sold in both the North and the South, and I assume transport costs are negligible.[15]

At each point during a product's lifetime, the firm makes a joint production and location

[12] In the empirical implementation of the model, I will include sector fixed effects that absorb M_j as well as any other fixed characteristics that differ across sectors.

[13] This is essentially an assumption that the North has a significant comparative advantage in managing innovation and R&D (see Antràs 2005 and Antràs and Helpman 2004), which is in line with evidence that the vast majority of worldwide R&D and patenting take place within OECD countries (OECD 2004). Imitators will not face this constraint, however, because they need not perform original innovation.

[14] Because the model provides a theoretical motivation for how international investment will respond to patent reform events in the medium term, I do not include longer-run considerations such as endogenous increases in Southern wages and Northern innovation rates. These general equilibrium effects are of clear theoretical and practical importance (Grossman and Helpman 1991) viewed from an aggregate perspective, but we lack sector-specific evidence detailing their influence across industries. I do establish in section 6 that the empirical results are robust to including country-year fixed effects, however, which would absorb any such changes under the model assumption that all industries face a common wage.

[15] In the Appendix, I consider an extended version of the model that includes both transport costs and fixed costs of establishing a foreign subsidiary. I find similar cross-sector predictions using either approach.

decision. As a monopolist, the firm maximizes profit by solving $\max_{x_h, x_l} pq - w^N x_h - w^i x_l$ when manufacturing in country i. Optimality implies that firms manufacturing in the North will charge prices p^N and earn per-period revenues r^N and profit π^N in each country as follows

$$p^N = \frac{w^N}{\rho}, \qquad r^N = \lambda \left(\frac{\rho}{w^N}\right)^{\sigma-1}, \qquad \pi^N = r^N(1-\rho). \tag{3}$$

Similarly, when products are manufactured in the South, prices are lower and per-period revenues and profit earned in each country are higher

$$p^S = \frac{(w^N)^\alpha (w^S)^{1-\alpha}}{\rho}, \qquad r^S = \lambda \left(\frac{\rho}{(w^N)^\alpha (w^S)^{1-\alpha}}\right)^{\sigma-1}, \qquad \pi^S = r^S(1-\rho). \tag{4}$$

Notice that π^i denotes the profit a firm makes in either market (North or South) when it locates production in country $i \in \{N, S\}$.

The innovator's monopoly power may be disrupted by imitation, however, and the risk of imitation is directly related to its manufacturing location. I assume manufacturing requires revealing proprietary information to assembly-line employees, and that this technology transfer in turn enables local entrepreneurs to more readily obtain access to the information.[16] Without access to this information, I assume potential imitators are unable to properly reverse-engineer patented products, preventing imitation even where intellectual property rights are weak. Cross-border access to this product-specific know-how is assumed to be prohibitively costly, so that imitators are constrained to pursue only those varieties that have been locally manufactured.

The risk of imitation affects innovators' sourcing decisions, because successful entry by an imitator may result in profit losses. Specifically, innovating firms competing with an imitation product capture only a fraction of the per-period profit described in (3) and (4). 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, as in 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:

[16] This proprietary information is distinct and complementary to that available by observing the final product and its associated patents, and can be accessed only by individuals outside the firm by interacting with employees familiar with the proprietary information or by observing the production facility. The assumption that acquiring proprietary information is less costly when it is in active local use is supported by recent evidence that former multinational employees are a significant conduit for technology transfer between multinational and domestic firms (Poole 2009).

$\xi_N = 1$ and $\xi_S < 1$. Only where a patent fails to be enforced may imitation products directly compete with innovators. Hence, imitation products may only be sold in the South.[17]

Endogenous Imitation

A fringe of potential imitators exists in both the North and the South. Any imitator with access to the proprietary information necessary for production can invest c to begin reverse-engineering a product. 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 m follows a uniform distribution over this period.[18] Imitation effort thus may or may not generate an imitation product within the targeted variety's lifetime.

When an imitator successfully enters a market with imperfect patent protection, it engages in Bertrand competition with the innovating firm wherever patents are not enforced, until the variety becomes obsolete (Grossman and Helpman 1991). Imitators produce final goods with the same production technology as innovating firms, but source headquarters services locally and thus have marginal costs w^S. This implies that only Southern entrepreneurs have an incentive to imitate products, because only they enjoy lower production costs than Northern multinationals. With this cost advantage, Southern entrepreneurs can profitably capture a market if patents are not enforced by charging a price just below the original innovator's marginal cost. In this scenario, successful imitators' prices, per-period revenues, and per-period profit will be

$$p_{im}^S = (w^N)^\alpha (w^S)^{1-\alpha}, \qquad r_{im}^S = \lambda \left((w^N)^\alpha (w^S)^{1-\alpha} \right)^{1-\sigma}, \qquad \pi_{im}^S = r_{im}^S \left(1 - \frac{w^S}{(w^N)^\alpha (w^S)^{1-\alpha}} \right). \tag{5}$$

If patents are enforced, but only imperfectly ($0 < \xi_S < 1$), imitators will charge p_{im}^S as above, but will earn only a fraction $(1 - \xi_S)$ of profit π_{im}^S.

By comparing imitators' expected profit with the cost of entry c, it is possible to determine which products will be selectively targeted by imitators, in turn influencing the sourcing decisions of Northern firms. Whether a potential imitator with access to the proprietary information will choose to pay c to pursue imitation of a particular variety of maturity t depends on a) the maximum per-period profit from imitation π_{im}^S, b) the variety's remaining

[17] See Grossman and Lai (2004) for further discussion.

[18] 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 product lifecycle lengths.

economic lifetime $\tau_j \equiv T_j - t$, and c) the strength of intellectual property institutions in both markets, ξ_N and ξ_S, defined above.

Assuming that there is no time discounting, imitators will have an incentive to reverse-engineer a sector-j product of maturity t whenever the net present value of doing so exceeds the fixed cost of imitation c. This will be true whenever the following inequality holds

$$(1-\xi_S)\pi_{im}^S E[\max\{0, T_j - t - m\}] = (1-\xi_S)\pi_{im}^S \frac{(T_j - t)^2}{2\overline{m}} > c. \tag{6}$$

Equation (6) shows that an imitator will earn a profit flow $(1-\xi_S)\pi_{im}^S$ for a time $T_j - t - m$ if reverse-engineering succeeds prior to obsolescence, that is if $m < T_j - t$.[19] The expression (6) implies that all products with at least

$$\overline{\tau}(\xi_S) = \left(\frac{2\overline{m}c}{\pi_{im}^S(1-\xi_S)}\right)^{1/2} \tag{7}$$

time remaining until obsolescence will face imitation risk if Southern imitators obtain the proprietary know-how necessary for production, which can occur once the innovating firm has offshored manufacturing to the South. On the other hand, it is apparent that:

Lemma 1: *Products in industries with $T_j < \overline{\tau}(\xi_S)$ do not face imitation risk. Longer-lived products that have only been manufactured in the North and are sufficiently mature ($t > T_j - \overline{\tau}(\xi_S)$) similarly face no imitation risk when production shifts to the South.*

2.3 The Product Cycle

Firms in each sector will make profit-maximizing production and location decisions, taking into account differences in intellectual property institutions and associated imitation risks across countries. But as shown in (6), imitation risk will not affect firms in all industries equally. Knowing this, a sector-j firm will select the optimal product maturity $t_j^* \in [0, T_j]$ at

[19] Because m is uniformly distributed on $[0, \overline{m}]$ with $\overline{m} > \max_j\{T_j\}$, the truncated expected value of $T_j - t - m$ is $\frac{(T_j - t)^2}{2\overline{m}}$.

which Southern production will begin by maximizing the following lifetime profit function

$$\Pi_j(t) = \begin{cases} 2\pi^N t + 2\pi^S E[\min\{T_j - t, m\}] \\ \qquad\qquad +(1+\xi_S)\pi^S E[\max\{0, T_j - t - m\}], & \text{if } t \leq T_j - \overline{\tau}(\xi_S) \\ 2\pi^N t + 2\pi^S (T_j - t), & \text{if } t > T_j - \overline{\tau}(\xi_S). \end{cases} \tag{8}$$

Equation (8) reveals the effect of imperfect patent enforcement in the South on Northern multinationals' sourcing incentives. The first case shows that relocating production to the South at maturity t when the product has a relatively large lifetime remaining ($T_j - t > \overline{\tau}(\xi_S)$) exposes the firm to imitation risk and the associated chance of profit reductions. Selling to both markets, the firm earns $2\pi^N$ with certainty until relocation at time t, but once manufacturing in the South begins, imitation occurs at an uncertain date $t + m$. The firm 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 will then be π^S only where patents are enforced, namely in the Northern market and in a fraction ξ_S of the Southern market; the resulting post-imitation profit is thus $(1+\xi_S)\pi^S$. In contrast, the second case in equation (8) shows that when relocation is postponed until a product has fewer than $\overline{\tau}(\xi_S)$ remaining years before obsolescence, the product does not face imitation risk. A direct implication of this is that fast-turnover products and relatively mature products face no imitation risk, and thus are sourced in the South where production costs are lower. I summarize this observation in Lemma 2:

Lemma 2: *Products in industries with $T_j < \overline{\tau}(\xi_S)$ are always manufactured in the South. In all other industries, production is shifted to the South at a maturity t no greater than $T_j - \overline{\tau}(\xi_S)$.*[20]

In (8), maximization of $\Pi_j(t)$ over possible relocation maturities t reveals a time-to-obsolescence sourcing cut-off $\tau^*(\xi_S) \equiv T_j - t_j^*$ that is invariant across sectors with different product lifecycle lengths. Varieties with less than $\tau^*(\xi_S)$ time remaining before obsolescence will be manufactured in the South, while all other varieties will be manufactured in the North. The value of this cut-off depends on ξ_S, innovators' π^N and π^S, and imitators' π^S_{im},

[20] As described at length in Vernon (1966) and more recently in Antràs (2005), there are reasons aside from imitation that might lead a firm to initially source production in the North and later in the South. For example, the headquarters-intensity of production may decline over the product lifecycle as in Antràs (2005). This force is orthogonal to the imitation-based mechanism I describe.

m, and c. Specifically, it can be shown that innovating firms will produce in the South only at maturities high enough to avoid imitation risk $\tau^*(\xi_S) = \overline{\tau}(\xi_S)$ when $(\frac{\pi^S - \pi^N}{\pi^S})^2 \leq (1-\xi_S)\frac{c}{2\overline{m}\pi^S_{im}}$.[21] Intuitively, this condition holds when the profit advantage of manufacturing in the South $(\frac{\pi^S - \pi^N}{\pi^S})$ is low relative to profit losses caused by imitation $(1-\xi_S)$. On the other hand, if the advantage of Southern production is high compared with losses from imitation, that is, if $(\frac{\pi^S - \pi^N}{\pi^S})^2 > (1-\xi_S)\frac{c}{2\overline{m}\pi^S_{im}}$, firms will also manufacture less-mature products in the South: $\tau^*(\xi_S) = \frac{\pi^S - \pi^N}{\pi^S(1-\xi_S)}2\overline{m} > \overline{\tau}(\xi_S)$. I summarize these results in the following equation, which shows that the sourcing cut-off is

$$\tau^*(\xi_S) = \begin{cases} \overline{\tau}(\xi_S), & \text{if } (\frac{\pi^S - \pi^N}{\pi^S})^2 \leq (1-\xi_S)\frac{c}{2\overline{m}\pi^S_{im}} \\ \frac{\pi^S - \pi^N}{\pi^S(1-\xi_S)}2\overline{m}, & \text{if } (\frac{\pi^S - \pi^N}{\pi^S})^2 > (1-\xi_S)\frac{c}{2\overline{m}\pi^S_{im}}, \end{cases} \tag{9}$$

where $\overline{\tau}(\xi_S)$ is defined in (7). It can be shown using (9) that $\frac{d\tau^*(\xi_S)}{d\xi_S} > 0$, hence the sourcing cut-off is increasing in the strength of Southern patent protection ξ_S. Notice that only in the bottom case described in equation (9) will imitation occur in equilibrium: $\tau^*(\xi_S) > \overline{\tau}(\xi_S)$. I focus on this case in the remaining analysis.[22]

Equation (9) implies that products will be manufactured in the North for $\max\{0, T_j - \tau^*(\xi_S)\}$ time, and in the South for $\max\{T_j, \tau^*(\xi_S)\}$ time. Because the age distribution of products within a sector is uniform with density $\psi_j(t) = 1/T_j$, this further implies that the measure of products manufactured in the North is

$$N_j(\xi_S) \equiv \int_0^{\max\{0, T_j - \tau^*(\xi_S)\}} \psi_j(t)dt = \int_0^{\max\{0, T_j - \tau^*(\xi_S)\}} \frac{1}{T_j}dt = \max\{0, \frac{T_j - \tau^*(\xi_S)}{T_j}\},$$

which is weakly increasing in T_j.[23] I summarize this result in Proposition 1:

Proposition 1: *The measure of sector-j products manufactured in the North $N_j(\xi_S)$ is (weakly) increasing in product lifecycle lengths T_j. $N_j(\xi_S)$ is zero for industries with $T_j < \tau^*(\xi_S)$, and is strictly increasing in T_j for all other industries.*

[21] These results are based on the maximization of (8) with respect to t. Because $\overline{\tau}(\xi_S)$ is a lower bound for $\tau^*(\xi_S)$, the optimal cut-off is either $\tau^*(\xi_S) = \overline{\tau}(\xi_S)$ or bigger. The optimum of the first case in (8) is $\frac{\pi^S - \pi^N}{\pi^S(1-\xi_S)}2\overline{m}$, which exceeds $\overline{\tau}(\xi_S)$ when $(\frac{\pi^S - \pi^N}{\pi^S})^2 > (1-\xi_S)\frac{c}{2\overline{m}\pi^S_{im}}$.

[22] The existence of imitation in equilibrium relies, in part, on my earlier assumption that m is uncertain to both imitators and innovators.

[23] Note that because T_j is fixed in each sector, the dependence of $N_j(\xi_S)$ on T_j is simply reflected by the subscript j on $N_j(\xi_S)$.

The steady-state distribution of $N_j(\xi_S)$ across sectors is illustrated in the top panel of Figure 3. I plot $N_j(\xi_S)$ as a function of sectors' product lifecycle lengths T_j at two different levels of Southern patent protection, ξ_S (solid) and ξ'_S (dashed), with $\xi_S < \xi'_S$.

In expectation, Southern affiliates in sector-j earn aggregate revenues $R_j(\xi_S)$ as follows, obtained by integrating variety-specific revenues over the distribution of product maturities

$$R_j(\xi_S) = \int_{\max\{0, T_j - \tau^*(\xi_S)\}}^{T_j} \left(2r^S(1-\kappa_{im}(t)) + (1+\xi_S)r^S\kappa_{im}(t)\right)\psi_j(t)dt, \tag{10}$$

where $\psi_j(t) = 1/T_j$ is the density of product maturities and $\kappa_{im}(t)$ is the probability that a maturity-t product has been imitated. After some simplification, the expression above can be reduced to

$$R_j(\xi_S) = \begin{cases} 2r^S, & T_j < \overline{\tau}(\xi_S) \\ 2r^S \cdot \left(1 - \frac{T_j}{2\overline{m}}\right) + (1+\xi_S)\cdot r^S \cdot \frac{T_j}{2\overline{m}}, & T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi_S)] \\ 2r^S \cdot \left(\frac{\tau^*(\xi_S)}{T_j} - \frac{\tau^*(\xi_S)^2}{2\overline{m}T_j}\right) + (1+\xi_S)\cdot r^S \cdot \frac{\tau^*(\xi_S)^2}{2\overline{m}T_j}, & T_j > \tau^*(\xi_S). \end{cases} \tag{11}$$

The first case in (11) shows that Southern affiliates earn full revenues $2r^S$ in industries with the shortest product lifecycles, because they are never imitated. In the second case, products are always sourced in the South but face imitation risk; it is apparent that 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 third case, products have longer lifetimes and only a measure $\tau^*(\xi_S)/T_j$ 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 $(1+\xi_S)r^S$.

2.4 Response to Improved Intellectual Property Rights

Suppose the South enacts a policy change that improves local patent enforcement from ξ_S to $\xi'_S > \xi_S$. Increases in ξ_S reduce the product maturity at which manufacturing in the South optimally begins. This is evident by observing that $\tau^{*\prime}(\xi_S) > 0$ in equation (9).

A straightforward implication of Proposition 1 is therefore that the difference in the measure of varieties sourced in the North at ξ_S and at $\xi'_S > \xi_S$ depends on T_j as follows

$$N_j(\xi_S) - N_j(\xi'_S) = \begin{cases} 0, & T_j < \tau^*(\xi_S) \\ \frac{T_j - \tau^*(\xi_S)}{T_j}, & T_j \in [\tau^*(\xi_S), \tau^*(\xi'_S)] \\ \frac{\tau^*(\xi'_S) - \tau^*(\xi_S)}{T_j}, & T_j > \tau^*(\xi'_S). \end{cases} \tag{12}$$

I plot $N_j(\xi_S) - N_j(\xi'_S)$ in the middle panel of Figure 3. This extensive-margin effect can be interpreted as the measure (level) of varieties that is immediately offshored to the South following the patent reform.[24] In fast-turnover industries with product lifecycles shorter than the original sourcing cut-off ($T_j \leq \tau^*(\xi_S)$), firms will not respond to the reform; in these industries, varieties were already manufactured in the South for their full lifetime at ξ_S and will continue to be at ξ'_S. In industries with longer product lifecycle lengths $T_j > \tau^*(\xi_S)$, 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 for $T_j > \tau^*(\xi'_S)$. $N_j(\xi_S) - N_j(\xi'_S)$ is thus a non-monotonic function of T_j.[25] This result is summarized in Proposition 2:

Proposition 2: *The increase in the measure of sector-j varieties manufactured in the South $N_j(\xi_S) - N_j(\xi'_S)$ following a patent reform from ξ_S to ξ'_S is a non-monotonic 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 will occur in the industry with $T_j = \tau^*(\xi'_S)$.*

Revenues earned by Southern affiliates are impacted both because of newly-shifted manufacturing activity and because existing imitated varieties capture a larger share of Southern sales under the stronger patent regime. Building from equations (11) and (12), it is apparent that the change in industry-j revenues earned by Southern affiliates following a Southern

[24] Notice that this effect may also be interpreted as a change in the *fraction* of offshored varieties in sector j. My empirical approach, described in section 4 below, is consistent with either interpretation, because all regressions include sector fixed effects that absorb the overall size of each industry.

[25] The non-monotonicity prediction of the model also holds under weaker assumptions. In particular, when I introduce random noise around the sourcing cut-off $\tau^*(\xi_S)$ so that it varies across sectors, I still find that firms in the shortest-lifecycle sectors do not respond to patent reforms, while firms in longer-lifecycle sectors on average respond by expanding activity, with intermediate sectors exhibiting the largest expansion. However, at upper values of T, the predicted shape of the response (as a function of T_j) depends on the underlying distribution of product maturities within each sector. The true shape across high-T sectors is therefore an empirical question that I resolve in my analysis to follow.

patent reform from ξ_S to ξ'_S depends on T_j

$$R_j(\xi'_S) - R_j(\xi_S) = \begin{cases} 0, & T_j < \overline{\tau}(\xi_S) \\ r^S \cdot (\xi'_S - \xi_S) \cdot \frac{T_j}{2\overline{m}}, & T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi_S)] \\ r^S \cdot \left(2\frac{T_j - \tau^*(\xi_S)}{T_j} + (\xi'_S - \xi_S) \cdot \frac{[\tau^*(\xi_S)]^2}{2\overline{m}T_j}\right), & T_j \in [\tau^*(\xi_S), \tau^*(\xi'_S)] \\ r^S \cdot \left(2\frac{\tau^*(\xi'_S) - \tau^*(\xi_S)}{T_j} + (\xi'_S - \xi_S) \cdot \frac{[\tau^*(\xi_S)]^2}{2\overline{m}T_j}\right), & T_j > \tau^*(\xi'_S). \end{cases} \tag{13}$$

The top line of (13) shows that Southern affiliates in the fastest-paced sectors will not earn additional revenues following the patent reform. In these industries, there is no new entry or revenue because there is no imitation risk under the pre-reform policy ξ_S. The second line of (13) considers industries with product lifetimes between the imitation cut-off $\overline{\tau}(\xi_S)$ and the pre-reform sourcing cut-off $\tau^*(\xi_S)$. These products face imitation risk, but nevertheless are manufactured in the South at all maturities, thus the reform has no immediate effect for these products other than to raise the expected value of future profit. Additional revenues are thus accrued only by firms with previously imitated varieties, of which a fraction $\frac{T_j}{2\overline{m}}$ exist at any moment. Under the improved patent policy, these varieties enjoy patent protection in a larger share ξ'_S of the Southern market than before. The third line shows that for slightly longer-lived products, with lifecycle lengths between pre-reform and post-reform sourcing cut-offs, affiliate revenues will rise due to both new product entry and improved patent protection for already-imitated products; empirically, these separate effects will generate movement along extensive (number of affiliates) and intensive (size of existing affiliates) margins, respectively, but together impact industry-level revenues R_j as shown in (13). As described in equation (12), a measure $N_j(\xi_S) - N_j(\xi'_S) = \frac{T_j - \tau^*(\xi_S)}{T_j}$ of varieties in industry j will shift production to the South following the reform, where they will earn full revenues $2r^S$. In addition, the fraction of already-imitated varieties $\frac{[\tau^*(\xi_S)]^2}{2\overline{m}T_j}$ will earn higher revenues. The fourth line of (13) shows that the mass of entrants defined in (12) is bounded by the new sourcing cut-off $\tau^*(\xi'_S)$.

It can be shown that the change in sector-j revenues earned by Southern affiliates $R_j(\xi'_S) - R_j(\xi_S)$ is a non-monotonic function of T_j, increasing for $T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi'_S)]$ 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 3. The testable implication of this result is that the response to patent reforms, measured by the sector-level increase in affiliate sales in the South after the reform, will follow a non-monotonic function of sectors' product lifecycle lengths with the peak

impact on industries with mid-length product lifecycles. I summarize these points in the following proposition:

Proposition 3: *The change in revenues earned by Southern affiliates in sector-j following a patent reform from ξ_S to ξ'_S, $R_j(\xi'_S) - R_j(\xi_S)$, is a non-monotonic function of T_j. Specifically, it is zero for $T_j < \overline{\tau}(\xi_S)$, increasing for $T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi'_S)]$, and decreasing for $T_j > \tau^*(\xi'_S)$. The largest influence on sourcing will occur in the industry with $T_j = \tau^*(\xi'_S)$.*

2.5 Multiple Southern Countries

Intellectual property rights vary both within countries over time and across countries. Qualitatively identical predictions to those described above apply to comparisons in the cross-section between countries with different levels of patent protection. Suppose that there are two Southern host countries, each of size 1 with prevailing wage w^S, that are identical but for different patent institutions ξ_S and ξ'_S, with $\xi'_S > \xi_S$ as above. Proposition 1 applies to these countries collectively. In particular, it is simple to show that, similar to (12), the measure $N_j(\xi_S) - N_j(\xi'_S)$ of varieties manufactured only in the stronger-patent host country will be a non-monotonic 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 (13) and similarly non-monotonic in T_j. This leads to the following result:

Proposition 4: *The differences a) 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, $N_j(\xi_S) - N_j(\xi'_S)$, and b) between the revenues earned by Southern affiliates in sector-j 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 non-monotonic functions of T_j. Specifically, both differences are zero for $T_j < \overline{\tau}(\xi_S)$, increasing for $T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi'_S)]$, and decreasing for $T_j > \tau^*(\xi'_S)$. Both $N_j(\xi_S) - N_j(\xi'_S)$ and $R_j(\xi'_S) - R_j(\xi_S)$ are therefore largest in the industry with $T_j = \tau^*(\xi'_S)$.*

2.6 Firm Heterogeneity

In the presence of firm-heterogeneity, patent laws will have differential effects on sourcing decisions across firms with different productivity levels. Returning to the two-country

North-South world, I assume that productivity differences across innovators can be summarized by a positive firm-specific parameter $\varphi \in [\varphi_L, \varphi_H]$ that enters the production function multiplicatively as in (Melitz 2003), $q_k = A\varphi_k x_h^{\alpha} x_l^{1-\alpha}$, where A is defined as before. Firms with higher productivity draws face correspondingly lower marginal costs, resulting in higher revenues and profit whether manufacturing in the North

$$p_k^N = \frac{w^N}{\rho\varphi_k}, \qquad r_k^N = \lambda\left(\frac{\rho\varphi_k}{w^N}\right)^{\sigma-1}, \qquad \pi_k^N = r_k^N(1-\rho), \tag{14}$$

or in the South

$$p_k^S = \frac{(w^N)^{\alpha}(w^S)^{1-\alpha}}{\rho\varphi_k}, \qquad r_k^S = \lambda\left(\frac{\rho\varphi_k}{(w^N)^{\alpha}(w^S)^{1-\alpha}}\right)^{\sigma-1}, \qquad \pi_k^S = r_k^S(1-\rho). \tag{15}$$

Assume further that all imitators share a fixed productivity level (normalized to 1), but are able to infer the productivity of innovating firms by observing market prices.[26] Imitators combine this firm-level information with observed product maturities to advantageously target firms' products. In some cases, the productivity of the Northern innovator is high enough to offset the cost advantage of a Southern imitator, deterring imitation of the most-productive firms' products. Specifically, Southern entrepreneurs will have no incentive to target an innovating firm's product if they will be unable to undercut its price, thus a high level of productivity protects against imitation. This can be seen by evaluating imitators' prices, revenues, and profit following successful imitation of variety k, produced by a Northern firm with productivity φ_k

$$p_{im,\varphi_k}^S = \frac{(w^N)^{\alpha}(w^S)^{1-\alpha}}{\varphi_k}, \qquad r_{im,\varphi_k}^S = \lambda\left(\frac{(w^N)^{\alpha}(w^S)^{1-\alpha}}{\varphi_k}\right)^{1-\sigma},$$

$$\pi_{im,\varphi_k}^S = r_{im,\varphi_k}^S\left(1 - \frac{\varphi_k w^S}{(w^N)^{\alpha}(w^S)^{1-\alpha}}\right). \tag{16}$$

In (16), π_{im,φ_k}^S is only positive if $\varphi_k < \left(\frac{w^N}{w^S}\right)^{\alpha} \equiv \overline{\varphi}$, so that firms with productivity draws above $\overline{\varphi}$ will not be targeted by imitators at any maturity, regardless of Southern patent protection ξ_S. This implies that low-productivity firms will be more sensitive to Southern patent reforms than high-productivity firms will be. Restricting attention to firms with

[26] My assumption that Southern entrepreneurs do not inherit innovators' productivity through imitation is reasonable provided that the productivity of a multinational is composed of non-product characteristics such as management quality, corporate structure, marketing, sales networks, and so on.

$\varphi_k < \overline{\varphi}$, imitators' incentives to enter the Southern market are summarized by a time-to-obsolescence threshold similar to (7)

$$\overline{\tau}(\xi_S, \varphi_k) = \left(\frac{2\overline{m}c}{\pi^S_{im,\varphi_k}(1-\xi_S)} \right)^{1/2}, \quad \varphi_k < \overline{\varphi} \qquad \frac{\partial \overline{\tau}(\xi_S, \varphi_k)}{\partial \xi_S} > 0, \ \frac{\partial \overline{\tau}(\xi_S, \varphi_k)}{\partial \varphi_k} > 0 \quad (17)$$

that is increasing in both ξ_S and φ_k.[27] This implies that up to a point, high-productivity firms can manufacture less-mature products in the South than low-productivity firms, without facing additional imitation risk. The next proposition summarizes these implications of firm heterogeneity:

Proposition 5: *The differential effects of improved Southern patent protection are more pronounced for relatively unproductive firms, because high-productivity firms with $\varphi_k > \overline{\varphi}$ do not face imitation risk and therefore are not affected by the reform.*

3 Data and Descriptive Statistics

3.1 Product Lifecycle Lengths

Measuring product lifecycle lengths by industry is an important step toward testing the predictions of the model. To develop a suitable measure, it is helpful to first consider what is meant by a *product* in the context of the model. In the model, a multinational firm must transfer otherwise tacit proprietary information to its manufacturing facility for production to occur. By assumption, the durability of this proprietary information is precisely aligned with the market lifetime of the associated product: obsolescence strikes both the product and the proprietary information required to produce it simultaneously. Implicit in this view is a separability of proprietary information across product generations: the proprietary information required to produce each new product must be distinct from that required to produce any old product. This is important, because imitation will only be related to sourcing if successive product generations are technologically distinguished enough to prevent imitation of a new-generation product solely on the basis of an ability to imitate its obsolete predecessor.

[27] To be exact, $\overline{\tau}(\xi_S, \varphi_k)$ is increasing in φ_k whenever $\left(\frac{w^N}{w^S}\right)^\alpha > \rho/\varphi_k$. But since $\overline{\tau}(\xi_S, \varphi_k)$ only applies to firms with $\left(\frac{w^N}{w^S}\right)^\alpha > \varphi_k$, this is easily true assuming that $\varphi_k^2 > \rho$.

Viewed this way, the critical difference across industries is not the time lapse between versions of a product embodying a single innovative idea, but rather the durability of the innovative idea itself, which may span multiple versions of the same product. Thus, although direct measures of micro-level product creation and destruction are available for certain household goods sectors (Broda and Weinstein 2010), I focus instead on the rate of technology obsolescence as a proxy better-suited (from an imitation view) to the definition of product lifecycle lengths described here.

My measurement approach relies on detailed information in patent records documenting the timing of citations, whereby subsequent patents refer to an existing patent as relevant prior art. The time lag between a citation date and the grant date of the cited patent tells us when the patented information was relevant to a subsequent innovation. Similarly, the sector-level distribution of patent citation lags reveals information about the durability of patented technologies by industry, with long lags indicating that technologies exhibit lasting relevance to future innovation.[28] The duration of patent citations can therefore be used to capture the rate of technology obsolescence, which is what T is meant to reflect in my model. I use the simple average of each sector-specific citation lag distribution as my main proxy for product lifecycle lengths.[29,30] Further below, I also consider other moments of the citation lag distribution and an alternative measure of technology durability.

I develop my primary proxy for product lifecycle lengths using the NBER U.S. Patent Citations Data File (Hall et al 2001). Among other variables, the database records the application date of every citing patent for each cited patent filed between 1976–2006. By taking the difference between the application date of the citing patent and the grant date of the cited patent, I construct the distribution of citation lags for each patent. To preserve comparability across patents over time, I restrict attention to citations occurring within 15

[28]Long lags may also indicate technologies that are complex and difficult to reverse-engineer. I discuss the empirical implications of this possibility in section 4.2 below.

[29]The validity of this measure rests on the assumption that the durability of patented intellectual property is positively correlated with the durability of unpatented intellectual property.

[30]Previous work based on patent citation data has used the number of citations received by a patent as a proxy for the its value (Hall, Jaffe and Trajtenberg 2001, Harhoff et al 1999, and Trajtenberg 1990), while other work has attempted to capture the "generality" of patented innovations by measuring the range of patent classes from which its citations originate (Moser and Nicholas 2004). But to my knowledge, variation across sectors in the duration of patent citation lags has not been used before to measure the rate of technology obsolescence or product lifecycle lengths. In principle, patent renewal data offer an alternative source of information regarding the economic durability of a patented innovation; however studies based on this data have thus far focused on renewals as indications of patent value rather than the economic durability of the underlying technology. See, for example, Pakes and Simpson (1991), Schankerman (1998), and Bessen (2008).

years of the patent grant date for all patents.[31] I then compute the average citation lag by patent class, as defined by the cited patent:

$$\widehat{T}_j = \frac{1}{N_j} \sum_{k=1}^{N_j} \text{Average Lag}_k = \frac{1}{N_j} \sum_{k=1}^{N_j} \left[\frac{1}{n_k} \sum_{i=1}^{n_k} (t_i - t_k) \right]$$

In the expression above, Average Lag_k is the average forward citation lag for patent k, based on subsequent patents $i = 1 \ldots n_k$ that cite patent k as prior art within 15 years of k's grant date. A forward citation lag $t_i - t_k$ is the time lapse between patent i's application date and patent k's grant date. The sector-j proxy for product or technology obsolescence, $\widehat{T}_j$, is the mean Average Lag across all patents $k = 1 \ldots N_j$ belonging to class j. Notice that I place no restriction on the patent class of citing patents, to allow for the possibility that a patented technology may have relevance to downstream innovation not only within its own class, but also for other patent classes.

Mapping this patent-based proxy for product lifecycle lengths $\widehat{T}_j$ onto the BEA's industry classification involves using a publicly-available U.S. Patent and Trademark Office concordance between U.S. patent classes and Standard Industrial Classification (SIC) 3-digit industry codes.[32] In most cases, the mapping from patent classes to industry codes is not one-to-one or many-to-many, but is many-to-one: many patent classes correspond to a given 3-digit industry code. In these many-to-one cases, class-level measures are averaged across all patent classes matched to the SIC code in question.[33] My main empirical results in sections 5 and 6 are based on this industry-level measure.

There is considerable variation in this measure across industries as shown in Figure 4. In Table 2, I provide the industry names for the top five and bottom five industries, as well as the five industries around the 75th percentile, ranked by my measure of product lifecycle lengths. It is apparent that electronics, clockwork-operated devices, and computers

[31] I consider citations of all U.S. patents granted during 1976–1990, which allows me to observe a complete 15-year citation history for each cited patent. Grant dates of citing patents span 1976–2006.

[32] The USPTO's concordance is between the U.S. Patent Classification System, as of December 2005, and 41 unique product fields based on the 1972 SIC. The concordance can be downloaded from the USPTO website: ftp://ftp.uspto.gov/pub/taf/sic_conc/2005_diskette/

[33] For example, SIC 361 (Electric Transmission and Distribution Equipment), uniquely matches USPTO Class 191 (Electricity: Transmission to Vehicles), however SIC 287 (Agricultural Chemicals) matches eight separate patent classes: 568 (Organic Compounds), 514 (Drug, Bio-Affecting and Body-Treating Compositions), 435 (Chemistry), 504 (Plant-Protecting and Regulating Compositions), 564 (Organic Compounds), 71 (Fertilizers), 987 (Organic Compounds), and 424 (Drug, Bio-Affecting and Body-Treating Compositions). Some industries, for example those in the service sector, do not correspond to any patent classes, and are therefore omitted from the analysis.

have shorter product lifecycle lengths than machine products, shipping containers, and non-electric heating equipment. A natural question is whether sectors with short product lifecycle lengths are relatively innovation intensive; Table 2 shows that on average, fast-turnover sectors indeed have higher R&D intensities, suggesting that this measure reflects meaningful differences across sectors. I plot R&D intensity against T_j by sector in Figure 5, which reveals a negative correlation (-32%), but also indicates that these two sector-level measures capture distinct industry characteristics.[34] Correlations between $\widehat{T}_j$ and other key sector characteristics are as follows: capital-intensity, 2.7%; labor-intensity, -8.9%; patent effectiveness, 2.2%; secrecy effectiveness, -1.0%.[35]

Comparison with Other Measures of Technology Obsolescence

I compare my main proxy $\widehat{T}_j$ with several alternatives that I incorporate into robustness checks in section 6. First, I compare $\widehat{T}_j$ with related measures based on other moments of the citation lag distributions described above. In particular, pronounced skewness in patent values (Schankerman 1998; Allison and Lemley, 2004; Hall, Jaffe, and Trajtenberg, 2005) 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.[36] Accordingly, I construct alternative proxies for T based on the 75th and 85th percentiles of each citation lag distribution. At the patent class level, I find a high degree of correlation between these measures and my main proxy $\widehat{T}_j$ (75th percentile: 93%; 85th percentile: 88%).

As a second step, I compare $\widehat{T}_j$ with 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 Schankerman 1998; Lanjouw and Schankerman 2004); this literature has also recognized differences across industries in technology decay rates implied by renewals (Schankerman 1998; Bessen 2008). My measure is broadly consistent with the estimates in Schankerman

[34] I also compare my measure of product lifecycle lengths with the micro-level product turnover measure developed in Broda and Weinstein (2010). Among the eight comparable sectors, the correlation between the two measures is approximately 75%.

[35] Capital intensity and labor intensity are based on reported firm-level asset values, employment, and sales from the BEA; patent and secrecy effectiveness are from Cohen, Nelson, and Walsh (2000).

[36] For example, Bessen (2008) finds that patents held by large firms are more valuable and of longer average duration than patents held by small firms, individuals, and non-profit organizations—even though large firms face higher patent renewal fees.

(1998), which finds that technology decay rates are significantly higher in electronics and mechanical equipment than in the chemical and pharmaceutical industries. Bessen (2008) also finds higher implied patent values in the chemical and pharmaceutical industries, compared with electronics and computers. However, direct comparisons between estimated decay rates and $\widehat{T}_j$ are difficult because these estimates are available for broad technology categories while my proxy $\widehat{T}_j$ is available for hundreds of patent classes. I therefore construct a simple alternative measure of average patent lifetimes by patent class using USPTO renewal data and compare this with $\widehat{T}_j$ (see the Appendix for details). I find an overall correlation of approximately 30% between this renewal-based measure and $\widehat{T}_j$, suggesting that citation lag and renewal data, though distinct, capture overlapping information regarding the durability of patented technology. Further detail on data construction and the comparison across alternative measures of T may be found in the Appendix.

3.2 U.S. Multinational Activity Abroad

I use detailed panel data on the global operations of U.S.-based multinational firms from the Bureau of Economic Analysis (BEA) Survey of U.S. Direct Investment Abroad. These data provide information on U.S. parent companies and each foreign affiliate on an annual basis.[37] 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.[38] Table 1 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.

To analyze the influence of patent laws on global sourcing patterns, I use disaggregated information on the sales, employment, physical assets, and R&D expenditures of multina-

[37] Any U.S. person having direct or indirect ownership or control of ten 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, affiliates with total assets, sales, or net income (or loss) below $3 million that did not own another affiliate were exempt from participating in the survey. Foreign affiliates are required to report separately unless they are in both the same country and three-digit industry. Each affiliate is considered to be incorporated where its physical assets are located.

[38] A key advantage of the BEA data is its nearly complete coverage; in a typical benchmark year, the survey accounts for over 99 percent of affiliate activity. In 1994, for example, participating affiliates accounted for 99.8 percent of total assets, 99.7 percent of total sales, and 99.9 percent of total U.S. FDI. This reflects the requirement of participation for every U.S. 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.

tional affiliates located in 92 countries.[39] Affiliate sales are reported separately by three-digit industry code, making it possible to categorize affiliate activity by primary industry. I also compare affiliate sales with U.S. exports by sector using data from the U.S. Census Bureau.[40]

3.3 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.[41] The index documents the strength of patent rights in five distinct categories: 1) extent of coverage, 2) membership in international patent agreements, 3) provisions for loss of patent protection, 4) enforcement mechanisms, and 5) duration of protection. Each category is given a score between zero and one based on whether prevailing patent laws meet specific, objective criteria.[42] The overall index is the unweighted sum of these five sub-indexes, and thus ranges between zero to five, with higher values indicating stronger protection. My main results are based on this overall index; I also evaluate the importance of each individual sub-index in alternative specifications. A key feature of these indexes for this analysis is their availability during 1980–2005 for 122 countries, enabling the use of an inclusive sample of host countries and survey years. Note that because index values are available in five-year intervals, I match the year of each benchmark survey to the closest available index year. Summary statistics appear in Table 1.

4 Econometric Framework

The model presented in section 2 features specific predictions regarding the level and composition of multinational activity across countries with varied levels of intellectual property protection, and across sectors with different product lifecycle lengths. In this section, I describe the empirical approach used to test these predictions.

[39] Countries were included in the data set if a) any U.S. FDI was recorded in any of the benchmark years and b) the patent rights index described below was available for the host country in at least two periods.

[40] 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/.

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

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

4.1 Baseline Estimating Equation

Propositions 2–4 suggest that firms' sourcing decisions respond to host-country patent laws only in sectors with relatively long product lifecycle lengths. Among these industries, those with intermediate product lifecycle lengths are predicted to be the most sensitive to patent rights. Firms in industries with short product lifecycles are insensitive, however, because the rapid turnover of products and ideas in these sectors reduces exposure to imitation risk. To test these predictions, I estimate the following generalized difference-in-differences specification

$$MNC_{ijt} = \alpha + \beta \cdot IPR_{it} + \gamma_1 \cdot IPR_{it} \times T_j + \gamma_2 \cdot IPR_{it} \times T_j^2 + \eta_i + \eta_j + \eta_t + X_{it} + \epsilon_{ijt}, \quad (18)$$

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 lifecycle length of sector j. The main coefficients of interest γ_1 and γ_2 jointly capture the differential influence of patent laws on affiliate activity across product lifecycle lengths T_j. The model predicts that $\gamma_1 > 0$ and $\gamma_2 < 0$, with magnitudes consistent with the non-monotonicity prediction, as illustrated in the middle and lower panels of Figure 3. Specifically, the combined effect $\gamma_1 \cdot T_j + \gamma_2 \cdot T_j^2$ should be non-negative across the range of T_j in the data (6.5 to 11 years), and should reach a peak somewhere in the middle of this range, that is $T_{peak} = \frac{-\gamma_1}{2\gamma_2} \in [6.5, 11]$.[43]

I estimate the baseline specification (18) with several measures of affiliate activity MNC_{ijt}. To test the affiliate revenue predictions of Propositions 3 and 4, I define MNC_{ijt} to be affiliate sales at the country-industry-year level. I also provide estimates using two other measures, affiliate assets and employment, both proportional to affiliate revenue in the model. These are of interest because without observing prices or quantities, it is unclear how observed increases in affiliate sales map to actual affiliate expansion. The value of assets and employment, however, are more likely to vary only with changes in affiliates' actual output. In particular, estimates based on employment have a clean interpretation, because employment is recorded as the number of full-time workers rather than a nominal value.

Finally, to decompose these industry-level effects along the extensive and intensive mar-

[43] As described in Propositions 2–4, the peak impact of a Southern patent reform (or the cross-section analog) occurs in the sector that has product lifecycle lengths equal to the new sourcing cut-off $\tau^*(\xi'_S)$. This cut-off is implicitly constrained in (8) to be less than $\max_j\{T_j\}$, and cannot fall below the smallest T without contradicting the model's main prediction that short-lifecycle sectors will be relatively insensitive to host-country patent laws.

gins, I estimate (18) using the number of affiliates by sector (extensive margin) or sales, assets, and employment at the individual affiliate level (intensive margin). These two outcome variables are both related to the same measure-of-varieties prediction of Proposition 2. Theoretically, the response described in Proposition 2 could result in effects along two different extensive margins: (1) the number of firms with a Southern affiliate, per sector at any point in time, and (2) the number of varieties per firm that are manufactured in the South at any point in time. Both are of interest, however the data are not precise enough at the product level to evaluate the latter extensive margin directly. In the empirical implementation of the model, I will therefore look at (1) directly, and will look at (2) indirectly by evaluating the size of individual affiliates.

The baseline specification includes a number of important controls. Industry fixed effects η_j absorb omitted sector-specific characteristics, including differences in imitation costs c and timing $\overline{m}$, per-period profit π^N and π^S, factor intensities of production including α, the average productivity of firms, total industry size M_j, and the initial demand distribution across countries. Industry fixed effects also subsume the main effect of T_j. Country fixed-effects η_i control for fixed differences across countries, such as geography and legal origin, that may influence the level of local activity by U.S. multinationals. I also add a vector of time-varying host-country covariates X_{it} that includes the log of GDP per capita and a measure of the prevailing corporate tax rate;[44] these help to control for economic development and the general business climate, respectively. Year fixed effects η_t absorb changes over time that affect observed affiliate activity in all countries and industries equally, such as shipping costs, communication costs, and general macroeconomic conditions. The error term ϵ_{ijt} combines any omitted factors that affect patterns of foreign direct investment. 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.

4.2 Identification

Identification of γ_1 and γ_2 is based on variation in the strength of patent protection within countries over time, and variation in product lifecycle lengths across sectors. A key advan-

[44] As in other recent studies using the BEA data, I construct the relevant corporate tax rate directly from observed tax payments by affiliates to host-country governments. The measure is the ratio of tax payments to affiliate income, averaged across firms. See Branstetter, Fisman, and Foley (2006) and Branstetter et al (2010).

tage of introducing this latter variation across industries 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 changes that may also influence the location of multinational activity.[45] Upon joining the WTO, for example, countries align intellectual property protection with the minimum standards set forth in TRIPS, but also make substantial concurrent changes to trade policy and investment laws. Because the main effect β will tend to reflect multinationals' responsiveness not only to the strength of local patent laws, but also to the quality of general institutions and overall investment conditions, its interpretation is unclear; the magnitude of β cannot be attributed to the distinct influence of intellectual property laws.

The interpretation of γ_1 and γ_2, on the other hand, is not subject to this concern, because as argued in section 2, variation in product lifecycle lengths determines multinationals' sensitivity to formal patent laws and less so to general institutions or development. Firms' sensitivity to general property rights protection, for example, or the quality of financial institutions is theoretically independent of the durability of ideas required for production.[46] Cross-industry variation in product lifecycle lengths thus allows me to capture the effect of patent laws separately from the effects of general institutions, even for cases in which both improve simultaneously. For this reason, I will emphasize the differential effects captured by γ_1 and γ_2 as I interpret the results. Also, note that this identification strategy does not require identical T_j values in each country, although it is important that the ordering of industries remains relatively stable across countries. This is relevant if the product lifecycle were to itself 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. Yet, because T_j is measured with U.S. data and applied to all countries and years uniformly across the sample, this possible form of endogeneity in T_j with respect to patent laws is unlikely to be an empirical concern.

A second factor that could influence the results is the possibility that multinational activity may precede and in some cases contribute to patent reforms in developing countries.[47] To affect the estimates γ_1 and γ_2, the intensity of lobbying would need to be systematically related to product lifecycle lengths T_j. In particular, firms would have to lobby host-country

[45] Prior work has established significant correlations between patent laws and general institutions, including GDP per capita and market openness (Acemoglu et al 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% in 1982, 67% in 2004). For further discussion, see Qian (2007).

[46] I evaluate this empirically in section 6 with standard measures of general institutional quality.

[47] See, for example, Branstetter, Fisman, and Foley (2006).

governments with intensities corresponding to product lifecycle lengths in the same non-monotonic cross-sectoral pattern predicted by the model. In a sense, this would confirm the theory that sensitivity to patent protection follows a non-monotonic function of T_j, thus the interpretation of γ_1 and γ_2 would remain intact. However, the magnitude of these estimates could be influenced upward or downward, depending on the mechanisms underlying the lobbying interaction.[48]

Finally, we may be concerned that T_j reflects variation across sectors in the ease of imitation. For example, long product lifecycles could be indicative of barriers to imitation such as product complexity; it may be that firms in short-lifecycle sectors innovate with greater intensity because their products are simpler to imitate, and rapid innovation is the only way to survive. Similarly, if products are well-protected by patents, incumbent monopolists may have a lower incentive to innovate relative to the case in which patents provide ineffective protection. Although neither of these possibilities can be ruled out entirely, the first is likely to work against finding confirmation of the model's predictions, 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-lifecycle sectors would be less, not more, sensitive to patent laws than fast-turnover sectors—the opposite of the main theoretical results. On the second point, the data show that patent effectiveness and product lifecycle lengths are not systematically related, thus it is unlikely that long product lifecycle lengths result from barriers to imitation created by exceptionally effective patents.[49]

5 Main Results

5.1 Patent Laws and Affiliate Presence at the Industry Level

To evaluate the influence of patent laws on the broad spatial and sectoral pattern of foreign direct investment, I first estimate equation (18) with a limited dependent variable. I define MNC_{ijt} to be an indicator variable equal to one if positive affiliate sales are observed in country i and sector j during period t. This zero-positive margin is active in the data, and

[48]It is not clear whether estimates of γ_1 and γ_2 would be overstated or understated. If lobbying firms are made aware of patent reforms well in advance, these firms may respond by building capacity in the reforming host-country prior to the actual reform. In this case, differential responsiveness to observed reforms would appear weaker, generating downward bias in one or both coefficients γ_1 and γ_2. If lobbying firms do not receive advance notification of reforms, correlation between affiliate activity and patent reforms could instead inflate estimates of γ_1 and γ_2. I include several tests of anticipation in section 6.

[49]The correlation between product lifecycle lengths and a standard measure of patent effectiveness from Cohen, Nelson, and Walsh (2000) is approximately 2%.

therefore informative—for many country-industry pairs, there is no affiliate activity in at least one benchmark year. In the Appendix, I analyze this margin theoretically by adding a proximity-concentration trade-off to the basic framework developed in section 2 (see Brainard 1997; Helpman, Melitz and Yeaple 2004). I show that in this setting, the resulting pattern of zeros inherits the basic characteristic of my model: the existence of zeros is insensitive to patent protection in short-lifecycle sectors, and is negatively related to patent strength in long-lifecycle sectors.[50]

In Table 3, I present the corresponding estimates based on the full sample of 92 countries, 37 industries, and five benchmark years. The results provide strong support for the theoretical prediction described above. In column 2, I find evidence that sectors with relatively long product lifecycles T are more responsive to the strength of host-country patent protection in a simple specification with the standard set of controls described in section 4—country, year, and industry fixed effects, corporate tax rates, and the log of per-capita GDP. In column 4, I add an interaction between GDP per capita and T to better isolate this effect from the influence of overall development. I find a smaller but significant estimate under this relatively conservative approach.

In line with the non-monotonicity predictions of Propositions 3 and 4, columns 3 and 5 suggest that the influence of patent laws may be highest in sectors with mid-length product lifecycles. Comparing the positive linear and negative quadratic interaction coefficients, these estimates imply that patent laws have the largest impact on an industry T between $T = 9.3$ years (column 3, 25th-percentile T) and ten years (column 5, 75th-percentile T). Column 5, again, includes interactions between per-capita GDP and T; these estimates are thus conservative relative to column 3.

In columns 6 and 7, I show that the results are robust to controlling for product complexity. I include separate interactions between the patent index and a sector-level measure of R&D intensity.[51] High levels of R&D intensity may reflect multiple industry characteristics, but product complexity is likely a key determinant. Because complex products may be costlier to innovate than imitate, firms in R&D-intensive sectors could be more dependent on patent protection to protect revenues. Consistent with this reasoning, estimates in columns

[50] In the theoretical framework developed in section 2 and extended in the Appendix, the mechanism generating action along the sector-level intensive and zero-positive margins does not yield interaction between these two margins, and I therefore test the model predictions for the each associated outcome variable separately. It is possible to develop a richer framework with selection (e.g. Helpman, Melitz and Rubenstein 2008), and this would be interesting to consider, though to affect my empirical results the magnitude of such effects would need to be correlated with T across sectors.

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

6 and 7 reveal that R&D-intensive sectors are more likely to locate affiliates in countries with strong patent laws. I explore the economic significance of R&D intensity relative to product lifecycle lengths in regressions below.

For general comparison, I also estimate the main effect of patent laws on the distribution of affiliate activity, omitting interactions with product lifecycle lengths. The reported coefficient on the patent index in column 1 is indistinguishable from zero, which contrasts with the significant coefficients on the interaction variables reported in columns 2–6. This finding reveals the potential limitations of an identification strategy that relies only on time-series variation in patent laws. By averaging the influence of patent laws across heterogeneous industries, the true impact of these laws may be overlooked despite being significant. Finally, although all regressions shown in Table 3 are estimated using OLS, the results are nearly identical when estimated with a probit approach; in addition, these results are robust to using alternative proxies for T, including 75th-percentile and 85th-percentile citation lags as described in section 3.[52]

5.2 Affiliate Activity at the Industry Level

Propositions 3 and 4 predict that the overall effect of patent laws is a non-monotonic function of product lifecycle lengths, with the largest influence occurring in sectors with mid-length product lifecycles. In Table 4, I evaluate this prediction using three measures of affiliate activity (sales, assets, and employment), and find strongly supportive results.[53]

Columns 1–3 present estimates of (18) based on affiliate sales at the country-industry-year level. All specifications include interactions between per-capita GDP and T, and are thus conservative. The results in column 2 trace out a concave function reaching its peak at $T = 9.9$ years, just below the 75th percentile of the product lifecycle length distribution. When I include separate interactions between IPR and R&D intensity in column 3, I find nearly identical results. Column 1, on the other hand, shows that the average linear interaction effect alone is not significant. I repeat these tests for affiliate assets in columns 4–6 and affiliate employment in columns 7–9, and find very similar results.

[52] Results available upon request.

[53] I also test the model predictions in the cross section of countries with different levels of patent protection. I construct a country-level average product lifecycle length based on the distribution of multinational activity across sectors for each year: $\overline{T}_i = \sum_j \phi_{ij} T_j$, where ϕ_{ij} is sector j's share of country i's overall affiliate sales. I find that the estimated relationship between $\overline{T}_i$ and IPR_i is 0.832 (standard error = .0443) in 1994, and a similarly significant relationship obtains in all other sample years. These results suggest that even in the cross-section, countries with stronger patent protection attract proportionally more multinational activity in sectors with longer product lifecycles.

Importantly, the effects in Table 4 are also economically significant. The estimates in column 2 suggest that sectors with intermediate lifecycle lengths (75th-percentile T) expand by 12 percentage points more, on average, than sectors with rapid product turnover (10th-percentile T) and 8.6 percentage points more than slow sectors (95th-percentile T) following a one standard deviation improvement in Southern patent protection. When I include interactions with R&D intensity, I find product lifecycle lengths to be more economically significant than R&D intensity by a factor of four. Estimates in column 3 imply a differential response across T of 17 percentage points, compared to 4.4 percentage points between the 10th percentile and median industries by R&D intensity. Similar results obtain for assets and employment. Column 6 reveals a 20 percentage point differential increase in affiliate assets across T, compared with a 7.5 percentage point differential increase across R&D intensity. Comparative statics for Column 9 suggest comparable effects of 17 percentage points across T and 4.2 percentage points across R&D Intensity.

First-Differences

While the prediction of Proposition 4 concerns the pattern of multinational activity across countries with varied patent regimes, Proposition 3 provides a qualitatively similar prediction for how affiliate activity will change in response to a patent reform within a country. To test this latter prediction, I estimate a first-differenced version of (18)

$$\Delta MNC_{ijt} = \beta \cdot \Delta IPR_{it} + \gamma_1 \cdot \Delta IPR_{it} \times T_j + \gamma_2 \cdot \Delta IPR_{it} \times T_j^2 + \Delta\eta_t + \Delta X_{it} + \Delta\epsilon_{ijt}, \quad (19)$$

where I have defined ΔMNC_{ijt} to be an indicator for increased affiliate sales. Note that the constant term as well as country and sector fixed effects have dropped out of the regression equation.

Corresponding results appear in Table 5, and confirm the essential predictions of Proposition 3. Compared with the estimates in Table 4, columns 5 and 7 reveal a qualitatively similar pattern of sensitivity to reforms, although comparative statics suggest the economic significance of these results is relatively modest. Specifically, column 7 implies that a one-standard-deviation patent reform will raise the likelihood of increased affiliate sales by 3.1 percentage points more in the median-T sector versus the 10th-percentile sector.

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 columns 2, 3, and 6.

5.3 Patent Laws and the Number of Affiliates

In Table 6, I investigate the nature of the industry-level results shown above in greater detail. Specifically, I make use of the BEA's affiliate-level data to characterize the influence of patent laws on the number of firms with affiliates in a particular country, sector, and year. This is closely related to the predictions of Proposition 2 regarding the measure of offshored varieties by sector at different levels of patent protection. From a Southern country's viewpoint, this extensive-margin dimension of response is also of primary interest, because it is tied to potential welfare effects of patent reforms. An expansion in the number of varieties produced in the country following reform would suggest stronger patent laws do more than transfer revenues from imitators to existing multinationals; instead, industrial activity generated by new firms may offset the decline in local imitation, in addition to expanding technology transfer.

The results shown in Table 6 support the predictions of Proposition 2 across all specifications. These are based on equation (18) where I have defined the dependent variable to be the log number of firms with affiliates in a particular country, sector, and year. Whether or not I include interactions between GDP per capita and T, the estimates imply that stronger patent laws attract more affiliates in high-T sectors, with the largest effects found in mid-T sectors (columns 3, 5, and 7). Coefficients on the main interactions are all significant at the 1% level, and as in Table 4, also imply that T is economically significant. Estimates in column 5 suggest that a one-standard-deviation patent reform generates a 9.3 percentage point differential increase in the number of affiliates between the median and 10th-percentile industries. Column 7 suggests product lifecycle lengths are a more significant determinant of sensitivity to patent laws than R&D intensity—comparative statics are 15 percentage points for T versus 3.2 percentage points for R&D intensity.

5.4 Affiliate Size

To complete the decomposition of the industry-level results along the extensive and intensive margins, I next examine changes in the size of individual affiliates (intensive margin). The model offers three relevant insights here. First, existing affiliates may expand as patent laws improve if they sell a particular variety that has already been imitated; these affiliates are larger under strong patent regimes because intellectual property is protected in a greater

share of the market. Second, the forces described in Proposition 2 could increase the size of affiliates if, following host-country patent reforms, multi-product firms begin producing varieties through an affiliate earlier in the product lifecycle. Third, Proposition 5 suggests that less-productive multinationals will be more sensitive to patent laws, implying that expansion in the measure of varieties and revenues (Propositions 2–4) will occur within affiliates of smaller and less productive multinationals. Together, these effects suggest that affiliate-level expansion should be non-monotonic as in Proposition 2, but more pronounced among less-productive firms. Expansion may also be observed in long-lifecycle sectors for the small fraction of already-imitated firms. The overall effect of patent laws on the average size of affiliates within an industry is theoretically ambiguous, however: new entrants will be small, but since some existing affiliates expand, the overall average could be drawn up or down.

In Tables 7 and 8, I explore the effect of patent laws on affiliate size, as measured by sales. These estimates are based on an affiliate-level variant of (18)

$$MNC_{ijpt} = \alpha + \beta \cdot IPR_{it} + \gamma_1 \cdot IPR_{it} \times T_j + \gamma_2 \cdot IPR_{it} \times T_j^2 + \eta_i + \eta_j + \eta_t + X_{it} + X_{pt} + \epsilon_{ijpt}, \quad (20)$$

where MNC_{ijpt} represents the sales of parent p's country-i, sector-j affiliates during year t. The product lifecycle length T_j corresponds to the industry reported for the affiliate. I include parent sales and R&D expenditures in X_{pt} to control for differences across parent company operations.

The estimates reported in Table 7 underscore the ambiguity predicted by my theory and described above. In column 2, the negative interaction on $IPR \times T$ suggests a decline in average affiliate size, however the more conservative estimate in column 4 is insignificant. The non-monotonicity predictions, as tested in columns 3, 5, and 7, find only weakly significant support in the latter two specifications. Column 1, which as in other tables excludes interactions between IPR and T, shows that patent laws alone are not a significant determinant of affiliate sizes.

By contrast, Table 8 presents separate estimates for affiliates of high-productivity and low-productivity firms.[54] In each of the specifications in columns 1–9, it is apparent that patent laws have a highly significant differential influence on affiliate size, but only for less-productive firms. Affiliates of high-productivity firms appear unaffected by patent laws relative to affiliates of low-productivity firms, which is in line with the theoretical

[54] To evaluate firm productivity, I compute simple Solow residuals and assign firms with above-median residuals to the high productivity group.

predictions of section 2.6 that these firms have low marginal costs and thus face little competition by imitators even where patent laws are weak. The estimates indicate that affiliates of low-productivity parents in sectors with intermediate product lifecycles (median T) expand an average of between 2.4 and 6.1 percentage points more than in sectors with short product lifecycles (10th-percentile T) following a one standard deviation patent reform, while there is no average or differential response for affiliates of high-productivity parents. Interestingly, these effects are only present along the product lifecycle length dimension—when I omit interactions between intellectual property rights and product lifecycle lengths, I find that regardless of parent productivity patent laws are not able to explain affiliate size.[55]

5.5 Exports Versus Multinational Activity

A critical assumption of the model is that manufacturing activity requires the use of proprietary information, distinct from that available by observing the firm's final product. Were it instead possible to correctly reverse-engineer a product merely by observing it, patent laws would be irrelevant to the location of affiliate activity in *all* sectors, not just those with rapid technology cycles. Entrepreneurs would be able to obtain the product from any market and thereby replicate it without obtaining any additional information.

The assumed importance of this informational input also has implications for firms' decisions to serve a foreign market directly by an affiliate versus by exporting finished products at arms'-length.[56] In the model, firms produce in the North and export to the South until products are relatively mature. Once a variety reaches the critical time-to-obsolescence cut-off $\tau^*(\xi_S)$, Southern production begins, because the expected losses from imitation have fallen below the certain cost savings of manufacturing in the South. This cut-off is an increasing function of Southern patent laws ξ_S; stronger Southern patent laws thus imply that innovating firms sell to the South through an affiliate earlier in the product lifecycle. An improvement in Southern patent protection therefore tilts the balance of cross-border activity

[55] The results are qualitatively identical when the dependent variable is instead based on affiliate assets or employment. I also estimate first-differenced versions of (20) using an indicator for increased affiliate size. I find support for the non-monotonicity results (columns 3, 5, and 7) and strong evidence suggesting affiliates of low-productivity firms are significantly more likely to expand following patent reforms, whereas those of high-productivity firms are less sensitive. These results are available upon request.

[56] A third alternative is licensing to an unaffiliated party in the destination country. Data on licensing to unaffiliated parties is not recorded on a country-specific basis in the BEA data or U.S. export statistics, leaving unobserved a potentially significant channel of firm-level economic response to patent reforms. This precludes a direct test of the FDI-versus-licensing decision.

away from exports and toward sales by a local affiliate, with differential effects across sectors as predicted by Propositions 2–4. An alternative test of these propositions is thus based on the fraction of North-to-South sales accounted for by affiliate sales.

In Table 9, I provide estimates of (18) using $MNC_{ijt} = \frac{S_{ijt}}{S_{ijt}+X_{ijt}}$, the fraction of U.S. sales to country i in sector j during year t that are accounted for by affiliate activity, where S_{ijt} is affiliate sales and X_{ijt} is U.S. exports. The results echo the pattern of previous estimates—a small, but significant positive linear interaction in columns 2, 4, and 6, and a concave shape reaching its peak at an intermediate T in columns 3, 5, and 7. The estimates in column 7 reveal, as before, that cross-industry variation in T is approximately four times as economically significant as variation in R&D intensity, suggesting differential responses of 3.1 percentage points (T) and 0.9 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.[57]

6 Robustness Checks and Alternative Specifications

6.1 Robustness Checks

I subject the main results to a number of robustness checks to better establish their stability. First, to control for changes in affiliate activity due to omitted factors that may change non-linearly over time, I include a full set of country-year fixed effects in each regression; these fixed effects absorb factors that change over time within countries such as wages, demand, and competitiveness, as well as the main effects of patent laws, GDP per capita, and corporate tax rates. To adjust for the possibility that entire industries may undergo significant changes in size, competitiveness, or innovativeness across benchmark years, I also estimate all main specifications with sector-year fixed effects. 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 U.S. outward foreign direct investment. I perform additional tests based on other moments of the citation lag distribution (75th and 85th percentiles) as proxies for T to account for the possibility

[57]Note that although similar measures of affiliate sales relative to exports are standard (i.e. Antràs 2003, Nunn and Trefler 2008), they are only proxies. These measures are imperfect because they may not capture all modes of U.S. sales to country i. Also, U.S. export data contain outbound intrafirm sales that will not be reflected by the numerator S_{ijt}. In addition, it is possible that affiliates produce intermediate products that are ultimately incorporated into final goods in another sector, thus these measures may be noisy.

that the average may place a higher weight on unsuccessful innovations that receive very few citations. To control for the possibility that other sector characteristics may influence sensitivity to patent laws, I reevaluate all main results by including separate interactions with sector-level measures of labor intensity, capital intensity, and plant-level returns to scale.[58] Finally, to reduce the influence of measurement error in product lifecycle lengths, I cluster standard errors by sector. For each of these tests, I find qualitatively similar results.[59]

6.2 Flexible Estimation

To further evaluate the shape of cross-industry sensitivity to patent laws, I estimate specifications that allow coefficients to vary flexibly across the product lifecycle length distribution. Table 10 reports estimates based on an equation similar to (18) in which T is categorized by quartile (Q_1^T, Q_2^T, Q_3^T, and Q_4^T). A dummy corresponding to each of the top three quartiles is interacted with IPR_{it}, as follows

$$MNC_{ijt} = \alpha + \beta \cdot IPR_{it} + \sum_{k=2}^{4} \gamma_k \cdot IPR_{it} \times 1_{T_j \in Q_k^T} + \eta_i + \eta_j + \eta_t + X_{it} + \epsilon_{ijt}, \tag{21}$$

so that β captures the effect of IPR on fast-turnover sectors in the bottom quartile of the T distribution. An advantage of this approach is that the differential effect of patent laws, as reflected by the coefficients β, γ_2, γ_3, and γ_4, is unrestricted across quartiles. A consistent pattern emerges from the estimates in Table 10: the second and third quartile coefficients tend to be positive and larger than the first or fourth quartile coefficients, and the third-quartile estimates are significant across all specifications. This pattern indicates a rising and falling curve that peaks at an intermediate T, providing additional support for the non-monotonicity results in Tables 3–8.

This flexible approach is not limited to quartiles, and can be extended to include as many categories as distinct values for T appear in the data. Figures 6 and 7 illustrate the results of this maximum-flexibility approach, where I estimate an equation similar to (21),

[58] Following Helpman, Melitz, and Yeaple (2004), plant-level returns to scale is the average number of non-production workers per establishment based on data from the 1997 Census of Manufactures.

[59] I also perform tests using the renewal-based proxy for T described in section 3, and find that long-lifecycle sectors are significantly more sensitive to patent laws using this alternative measure. In addition, I use the within-industry variation in $\widehat{T}_j$ 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 3-digit sectors with below-median variation in $\widehat{T}_j$ across patent classes. I also perform tests based on the sub-components of the patent protection index, and find that enforcement and membership in international intellectual property treaties are the most important determinants of multinational activity.

but with a separate interaction between the patent index and a dummy for each sector. These coefficients are plotted against T in Figure 6 using the log number of affiliates as the dependent variable as in Table 6, and in Figure 7 using an indicator for any affiliate activity as in Table 3. In both cases, the data fit a concave curve that peaks at an intermediate T.[60] Together with the estimates in Table 10 and on the squared interaction terms appearing in the other tables, these results suggest that the model's non-monotonicity prediction is robust across a variety of estimation strategies.

6.3 Patent Effectiveness

Prominent 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. The effectiveness of these mechanisms is important to the present study because although product lifecycle lengths do not predict firms' reliance on patents *per se*, the model presumes that strong host-country patent protection will protect firms' patented innovations from imitation as advertised. The conclusions of the surveys mentioned above suggest that this is far from true, however, with certain sectors receiving little actual protection even in developed countries with strong patent laws. I incorporate this observation into my analysis using measures of patent- and secrecy effectiveness from Cohen, Nelson, and Walsh (2000).

In the top panel of Table 11, I present split-sample estimates categorized by whether a sector is above or below the median score for patent effectiveness, where high scores indicate high effectiveness. The emerging pattern of results shows that, as one would predict, the cross-industry mechanism of patent law sensitivity predicted by my model operates more forcefully in sectors for which patents are effective. I show two pairs of results for each of two dependent variables (log affiliate sales and log number of affiliates) and in all cases find significant coefficients only in high-effectiveness sectors.[61] These results also obtain using 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.[62] In unreported results, I take a similar approach

[60] Similar shapes obtain when estimation is based on intensive-margin variables such as sales at the industry or affiliate level.

[61] When patent effectiveness is included in specification (18) in place of T, the results indicate that this variable is also an important source of cross-industry variation in sensitivity to patent laws. Similarly, there is reasonable evidence that secrecy effectiveness is important as well.

[62] Results available upon request.

using secrecy effectiveness, and reach a similar conclusion: sectors for which secrets are ineffective tend to be more sensitive to host-country patent laws.

6.4 Reform Anticipation: TRIPS

If firms anticipate changes to patent laws, even indirectly, we may expect patterns of response to differ from the model predictions. Within the sample period, at least one major pre-announced reform took effect, providing an opportunity to measure the possible influence of anticipation on the results described above. Specifically, the phase-in of the WTO's TRIPS agreement created a lag between announcement and implementation dates for developing and less-developed countries. While the TRIPS agreement went into effect with the creation of the WTO on January 1, 1995, developing countries were permitted to delay meeting required minimum standards until 2000, while less-developing countries were given until 2006 (for pharmaceutical patents, this was extended to 2016). A large set of countries, identified in Kyle and McGahon (2009), thus implemented pre-announced reforms during the sample period.

In the middle panel of Table 11, I present estimates comparable to those in Tables 3, 4 and 6 based on a sample that excludes observations for countries permitted to delay implementation of the TRIPS reforms until 2000. I omit only the potentially affected years, 2000 and 2004. All specifications include interactions between GDP per capita and T, and columns 3, 6, and 9 also include interactions between IPR and R&D intensity. Compared with the full-sample estimates, no clear pattern of difference emerges, suggesting—if only indirectly—that policy anticipation is unlikely to have had a significant impact on my main results.

6.5 Placebo Test: Host-Country Financial Development

The model posits that patent laws will exhibit a strong differential effect on multinational activity depending on sectors' product lifecycle lengths T. Moreover, the identification strategy presumes that sensitivity to general host-country institutions will not be determined by product lifecycle lengths, so that the interaction variables $IPR \times T$ and $IPR \times T^2$ isolate the effects of patent laws separately from those of general institutions. In the lower panel of Table 11, I investigate the influence of one general institution, host-country credit conditions, on multinational activity as a placebo test. I re-run the main regressions using a measure of private credit conditions from Djankov, McLiesh, and Schleifer (2007) in place of IPR. The

coefficients on the key interactions are, as expected, insignificant.[63]

6.6 Growth in Affiliate Activity

Finally, although the model does not offer explicit predictions for the influence of patent laws on affiliate growth, I nevertheless investigate growth effects in Table 12. I apply the same empirical strategy and regression specifications as in previous tables, but here define the dependent variable to be the one-period change in the log of sector-level affiliate sales. The estimates in columns 2–7 suggest that patent laws have a significant differential effect on affiliate growth. As in previous tables, the main interaction between patent strength and product lifecycle lengths is positive; where a quadratic interaction is included, the estimates reveal sectors with mid-length lifecycles are the most sensitive to patent laws. These estimates indicate that the influence of patent laws on multinational activity extends beyond level effects, impacting growth. This further suggests that the complete effects of patent reforms on multinational activity may not be felt immediately, and may thus be larger and longer-lasting than those implied by my main results.

7 Conclusion

This paper has examined multinationals' sensitivity to host-country intellectual property institutions both theoretically and empirically. Within a model of firms' global sourcing decisions, I develop predictions for the spatial and sectoral composition of multinational activity. The model suggests that sensitivity to local patent institutions will be concentrated among sectors with relatively long product lifecycles, with the most pronounced sensitivity in sectors with mid-length product lifecycles. Among sensitive sectors, stronger host-country patent laws attract more aggregate affiliate activity as well as a larger number of affiliates. The size of the average affiliate also expands, but only among less productive firms.

These predictions find robust empirical support within a comprehensive panel of U.S. multinationals' activity spanning 92 countries and 37 industries during 1982–2004. Using the interaction between patent laws and product lifecycle lengths, I am able to explain systematic variation in affiliate activity, as measured by sales, assets, and employment, along all

[63] I also perform analogous placebo tests with measures of general institutional quality from Glaeser, et al (2004), which include constraints on the executive, autocracy, democracy, and years of schooling. As with private credit conditions, the results indicate that sensitivity to these general institutions does not vary systematically across sectors according to product lifecycle lengths.

margins of multinational activity. The results therefore provide evidence that cross-industry differences in the rate of product obsolescence are a significant determinant of firms' sensitivity to host-country patent laws. In addition, the results establish the causal influence of patent laws on multinational investment, which has challenged prior studies due to concurrent policy reforms.

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, 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 a key conduit for technology transfer, but primarily in the subset of sectors with relatively long-lived intellectual property.

Finally, it should be noted that my estimates may provide only a lower bound for the overall effect of host-country patent protection on industrial activity. This is because stronger patent protection may attract not only increased foreign direct investment, but also greater levels of arms'-length licensing. Stronger patent regimes may also encourage domestic innovation. Exploring these other channels of response, and measuring their magnitudes relative to impacts on domestic imitation and consumer surplus at different time horizons is an important area for future research.

References

Acemoglu, D., Johnson, S., and Robinson, J. (2005). "Institutions as a Fundamental Cause of Long-Run Growth," in Aghion, P. and Durlauf, S. (editors), *Handbook of Economic Growth*, vol. 1A, pp. 385–472.

Antràs, P. (2003). "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics*, vol. 118, no. 4, pp. 1375–1418.

Antràs, P. (2005). "Incomplete Contracts and the Product Cycle," *American Economic Review*, vol. 95, no. 4, pp. 1054–1073.

Antràs, P., Desai, M.A., and Foley, C.F. (2009). "Multinational Firms, FDI Flows, and Imperfect Capital Markets," *Quarterly Journal of Economics*, vol. 124, no. 3, pp. 1171–1219.

Antràs, P., Garicano, L., and Rossi-Hansberg, E. (2008). "Organizing Offshoring: Middle Managers and Communication Costs," in Helpman, E., Marin, D., and Verdier, T., (editors), *The Organization of Firms in a Global Economy*, Harvard University Press.

Antràs, P. and Helpman, E. (2004). "Global Sourcing," *Journal of Political Economy*, vol. 112, no. 3, pp. 552–580.

Arora, A., Ceccagnoli, M., and Cohen, W.M. (2003). "R&D and the Patent Premium," NBER Working Paper 9431.

Bernard, A., Jensen, J.B., Redding, S.J. and Schott, P. (2010). "Intra-Firm Trade and Product Contractibility," *American Economic Review Papers and Proceedings*, vol. 100, pp. 444–448.

Besedes, T. and Prusa, T. (2006a). "Ins, Outs, and the Duration of Trade," *Canadian Journal of Economics*, vol. 39, no. 1, pp. 266–295.

Besedes, T. and Prusa, T. (2006b). "Product Differentiation and Duration of U.S. Import Trade," *Journal of International Economics*, vol. 70, no. 2, pp. 339–58.

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

Bilir, L.K., Moser, P., and Talis, I. (2010). "Do Patent Treaties Encourage Technology Transfer? Evidence from the U.S. Accession to the Paris Convention in 1887," Stanford University mimeo.

Bils, M. and Klenow, P.J. (2004). "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, vol. 112, pp. 947–985.

Brainard, L. (1997). "An Empirical Assessment of the Proximity-Concentration Trade-off Between Multinational Sales and Trade," *American Economic Review*, vol. 87, no. 4, pp. 520–544.

Branstetter, L., Fisman, R.J., and Foley, C.F. (2006). "Do Stronger Intellectual Property Rights Increase International Technology Transfer? Empirical Evidence from U.S. Firm-level Panel Data," *Quarterly Journal of Economics*, vol. 121, no. 1, pp. 321–349.

Branstetter, L., Fisman, R.J., Foley, C.F., and Saggi, K. (2010). "Does Intellectual Property Rights Reform Spur Industrial Development?" *forthcoming, Journal of International Economics.*

Broda, C. and Weinstein, D.E. (2010). "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, vol. 100, no. 3, pp. 691–723.

Burk, D.L. and Lemley, M.A. (2009). *The Patent Crisis and How the Courts Can Solve It*, University of Chicago Press.

Chaudhuri, S., Goldberg, P., and Jia, P. (2006). "Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India," *American Economic Review*, vol. 96, no. 5, pp. 1477–1514.

Cohen, W.M., Nelson, R.R., and Walsh, J.P. (2000). "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (Or Not)," NBER Working Paper 7552.

Criscuolo, C., Haskel, J., and Slaughter, M. (2010). "Global Engagement and the Innovation Activities of Firms," *International Journal of Industrial Organization*, vol. 28, no. 2, pp. 191–202.

Davidson, W.H. and McFetridge, D.G. (1985). "Key Characteristics in the Choice of International Technology Transfer Mode," *Journal of Industrial Economics*, vol. 16, no. 2, pp. 5–21.

Djankov, S., McLiesh, C., and Shleifer, A. (2007). "Private Credit in 129 Countries," *Journal of Financial Econometrics*, vol. 84, no. 2, pp. 299–329.

Doms, M.E. and Jensen, J.B. (1998). "Comparing Wages, Skills, and Productivity between Domestically and Foreign-Owned Manufacturing Establishments in the United States," in Baldwin, R., Lipsey, R., Richards, D. (editors), *Geography and Ownership as Bases for Economic Accounting*, University of Chicago Press.

Evenson, R.E. (1990). "Survey of Empirical Studies," in Wolfgang E. Siebeck (editor), *Strengthening Protection of Intellectual Property in Developing Countries*, World Bank discussion paper no. 112.

Ginarte, J.C. and Park, W.G. (1997). "Determinants of Patent Rights: A Cross-National Study," *Research Policy*, vol. 26, no. 3, pp. 283–301.

Glaeser, E., La Porta, F., Lopez-de-Silanes, F., and Schleifer, A. (2004). "Do Institutions Cause Growth?" *Journal of Economic Growth*, vol. 9 no. 3, pp. 271–303.

Glass, A.J. and Saggi, K. (2002). "Licensing versus Direct Investment: Implications for Economic Growth," *Journal of International Economics*, vol. 56, no. 1, pp. 131–53.

Grossman, G.M. and Helpman, E. (1991). "Quality Ladders and Product Cycles," *Quarterly Journal of Economics*, vol. 106, no. 2, pp. 557–86.

Grossman, G.M. and Lai, E. (2004). "International Protection of Intellectual Property," *American Economic Review*, vol. 94, no. 5, pp. 1635–1653.

Hall, B.H., Jaffe, A.B., and Tratjenberg, M. (2001). "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools," NBER Working Paper 8498.

Harhoff, D., Narin, F., Scherer, F.M., and Vopel, K. (1999). "Citation Frequency and the Value of Patented Inventions," *Review of Economics and Statistics*, vol. 81, no. 3, pp. 511–515.

Harhoff, D. (2000). "R&D Spillovers, Technological Proximity, and Productivity Growth–Evidence from German Panel Data," *Schmalenbach Business Review*, vol. 52, no. 3, pp. 238–260.

Helpman, E. (1993). "Innovation, Imitation, and Intellectual Property Rights," *Econometrica*, vol. 61, no. 6, pp. 1247–1280.

Helpman, E., Melitz, M., and Rubenstein, Y. (2008). "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, vol. 123, pp. 441–487.

Helpman, E., Melitz, M., and Yeaple, S.R. (2004). "Export versus FDI with Heterogeneous Firms," *American Economic Review*, vol. 94, no. 1, pp. 300–316.

Heston, A., Summers, R., and Aten, B. (2009). Penn World Table Version 6.3, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

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

Krugman, P. (1979). "A Model of Innovation, Technology Transfer, and the World Distribution of Income," *Journal of Political Economy*, vol. 87, no. 2, pp. 253-66.

Kyle, M. and McGahon, A. (2009). "Investments in Pharmaceuticals Before and After TRIPS," NBER Working Paper 15468.

Lanjouw, J. (1993). "Patent Protection: Of What Value and for How Long?" NBER Working Paper 4475.

Lanjouw, J., Pakes, A., and Putnam, J. (1998). "How to Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data," *Journal of Industrial Economics*, vol. 46, no. 4, pp. 405–432.

Levin, R.C., Klevorick, A.K., Nelson, R.R., Winter, S.G., Gilbert, R., and Griliches, Z. (1987). "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity*, vol. 1987, no. 3, Special Issue On Microeconomics, pp. 783–831.

Lerner, J. (2009). "The Empirical Impact of Intellectual Property Rights on Innovation: Puzzles and Clues," *American Economic Review Papers and Proceedings*, vol. 99, pp. 343–348.

Manova, K.B. (2007). "Credit Constraints, Heterogeneous Firms, and International Trade," NBER Working Paper 14531.

Manova, K.B., Wei, S., and Zhang, Z. (2010). "Firm Exports and Multinational Activity under Credit Constraints," Stanford University mimeo.

Mansfield, E. and Schwartz, M. (1981). "Imitation Costs and Patents: An Empirical Study," *The Economic Journal*, vol. 91, no. 364, pp. 907–918.

McCalman, P. (2004). "Foreign Direct Investment and Intellectual Property Rights: Evidence from Hollywood's Global Distribution of Movies and Videos," *Journal of International Economics*, vol. 62, pp. 107–123.

Melitz, M. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, vol. 71, no. 6, pp. 1695–1725.

Moser, P. (2003). "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World Fairs," NBER Working Paper 9909.

Moser, P. and Nicholas, T. (2004). "Was Electricity a General Purpose Technology? Evidence from Historical Patent Citations," *American Economic Review Papers and Proceedings*, vol. 94, pp. 388–394.

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

Nunn, N. (2007). "Relationship-Specificity, Incomplete Contracts and the Pattern of Trade," *Quarterly Journal of Economics*, vol. 122, no. 2, pp. 569–600.

Nunn, N. and Trefler, D. (2008). "The Boundaries of the Multinational Firm: An Empirical Analysis," in Helpman, E., Marin, D., and Verdier, T., (editors). *The Organization of Firms in a Global Economy*, Harvard University Press.

Organization for Economic Cooperation and Development, (2004). "Patents and Innovation: Trends and Policy Challenges."

Pakes, A. (1986). "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, vol. 54, no. 4, pp. 755–784.

Pakes, A. and Simpson, M. (1991). "The Analysis of Patent Renewal Data," *Brookings Papers on Economic Activity, Microeconomic Annual*, pp. 331–401.

Park, W.G. (2008). "International Patent Protection: 1960–2005," *Research Policy*, vol. 37, no. 4, pp. 761–66.

Poole, J.P. (2009). "Knowledge Transfers from Multinational to Domestic Firms: Evidence from Worker Mobility," University of California, Santa Cruz mimeo.

Qian, Y. (2007). "Do National Patent Laws Stimulate Domestic Innovation In A Global Patenting Environment?" *Review of Economics and Statistics*, vol. 89, no. 3, pp. 436–453.

Schankerman, M.A. (1998). "How Valuable is Patent Protection? Estimates by Technology Field," *RAND Journal of Economics*, vol. 29, no. 1, pp. 77–107.

Schankerman, M.A. and Pakes, A. (1986). "Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period," *Economic Journal*, vol. 96, no. 384, pp. 1052–1077.

Trajtenberg, M. (1990). "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *Rand Journal of Economics*, vol. 21, no. 1, pp. 172–187.

UNCTAD (2005). *World Investment Report 2005: Transnational Corporations and the Internationalization of R&D*. United Nations, New York and Geneva.

Vernon, R. (1966). "International Investment and International Trade in the Product Cycle," *Quarterly Journal of Economics*, vol. 80, no. 2, pp. 190–207.

Yang, L. and Maskus, K.E. (2001). "Intellectual Property Rights, Licensing, and Innovation in an Endogenous Product-Cycle Model," *Journal of International Economics*, vol. 53, pp. 169–187.

Figure 1: Difference in the Composition of Affiliate Sales Across Countries, 1994

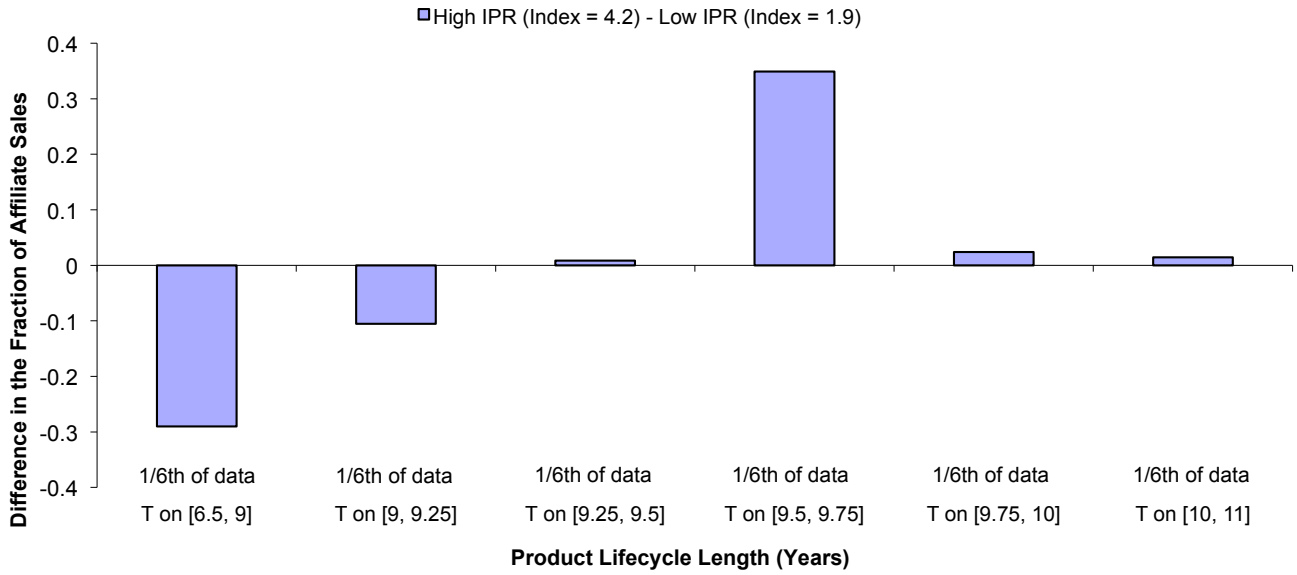

Notes: This figure is based on affiliate sales by industry as reported to the BEA in 1994. It shows the difference in the fraction of affiliate sales in strong-patent and weak-patent host countries, grouped by intervals according to product lifecycle lengths. Each bin represents 1/6th of the data. The six countries in each group are as follows. Weak patent: Guatemala, Indonesia, Malaysia, Paraguay, Thailand, Uruguay; Strong patent: Australia, Hungary, Japan, New Zealand, Norway, Sweden.

Figure 2: Change in the Composition of Affiliate Sales After Patent Reforms, 1994 - 2004

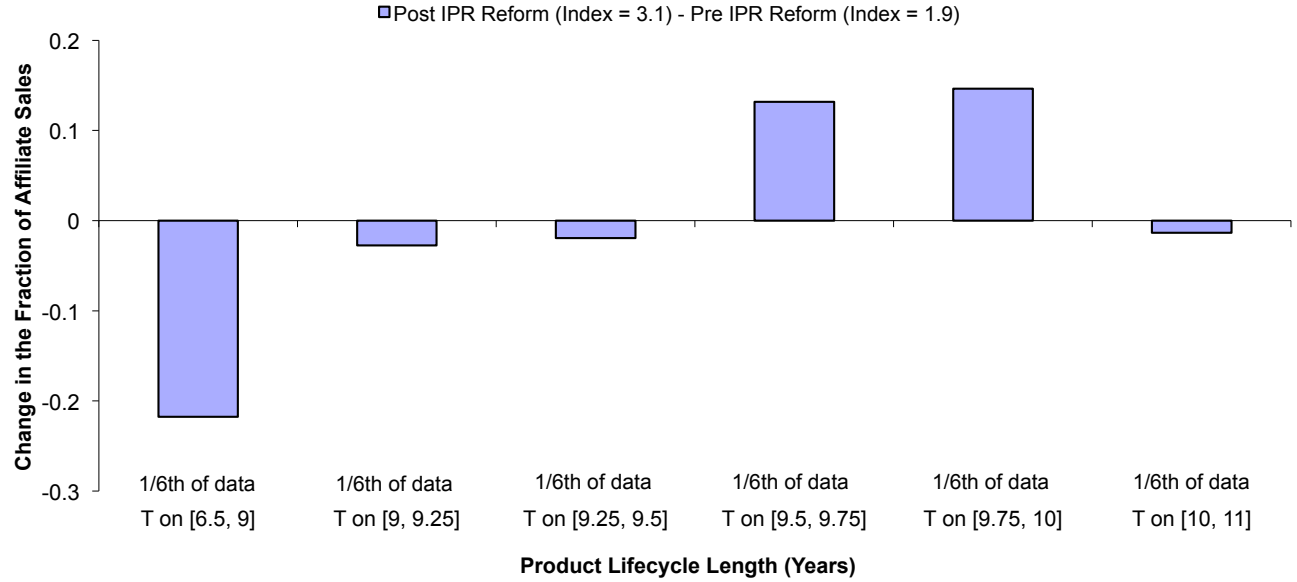

Notes: This figure is based on affiliate sales by industry as reported to the BEA in 1994 and 2004. For a group of six countries, the figure shows the difference in the fraction of affiliate sales during a period of strong patent protection (2004) versus a period of weak patent protection (1994). Each bin represents 1/6th of the data. Each country included (Guatemala, Indonesia, Malaysia, Paraguay, Thailand, Uruguay) implemented intellectual property rights reforms during 1994-2004.

Figure 3: Product Lifecycle Lengths and the Location of Manufacturing Activity

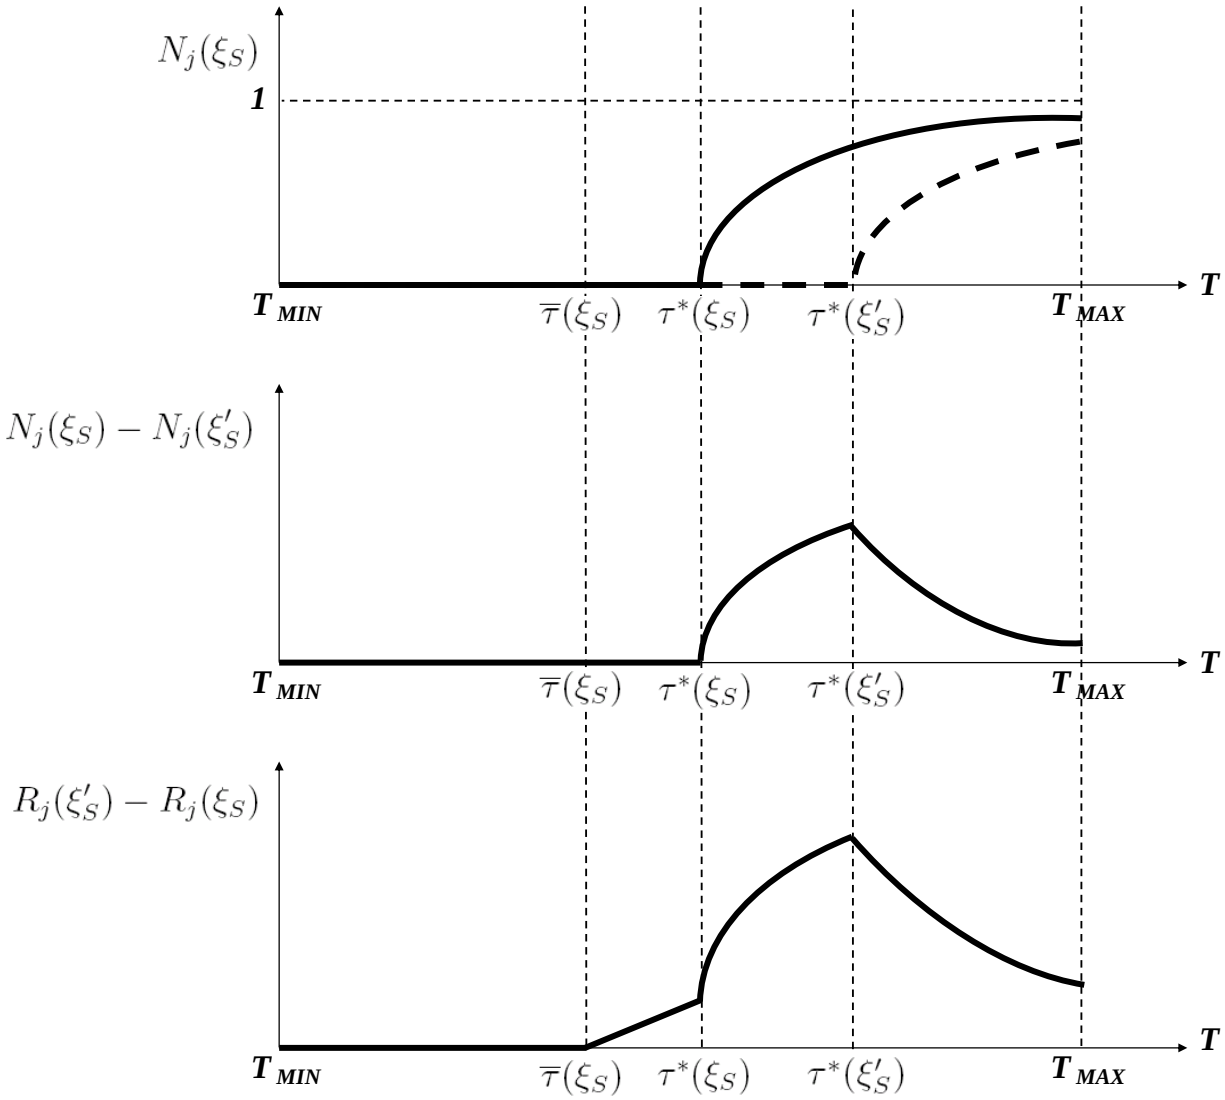

Notes: The upper panel of this figure shows the measure of product varieties manufactured in the North $N_j(\zeta_S)$ as a function of product lifecycle lengths T_j at two different levels of host-country patent protection ζ_S (solid) and ζ'_S (dashed), $\zeta_S < \zeta'_S$. The middle panel takes the 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 $T_j > \tau^*(\zeta_S)$. 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.

Figure 4: Average Patent Citation Lag, by Industry

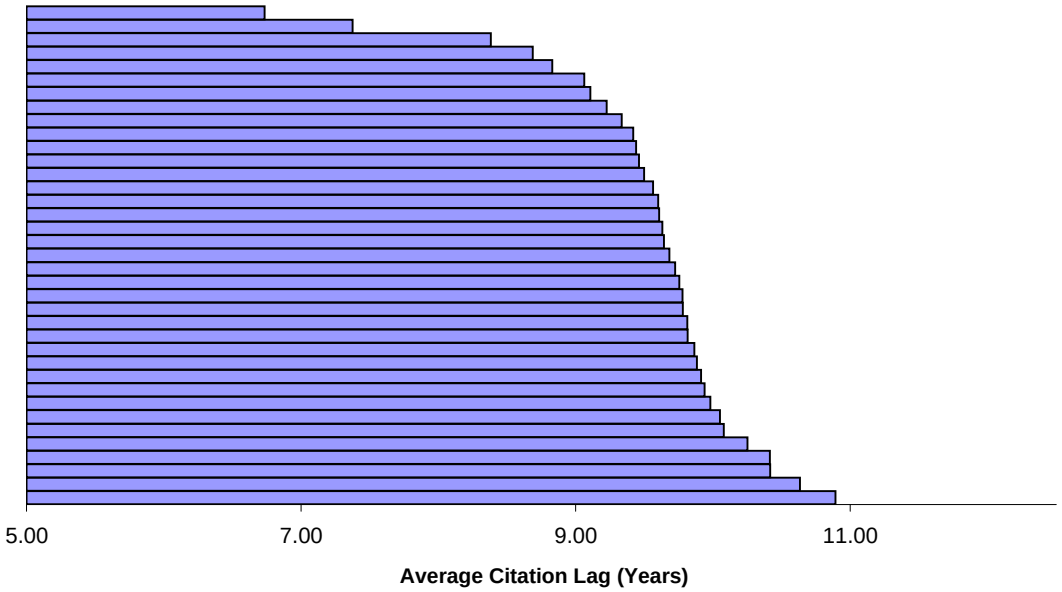

Notes: This figure shows the average patent citation lag for each of the 37 SIC 3-digit industries studied in this paper. Citation lags are measured in years and were computed using data in the NBER Patent Citation Datafile (Hall, et al 2001).

Figure 5: Product Lifecycle Lengths vs. R&D Intensities Across Industries

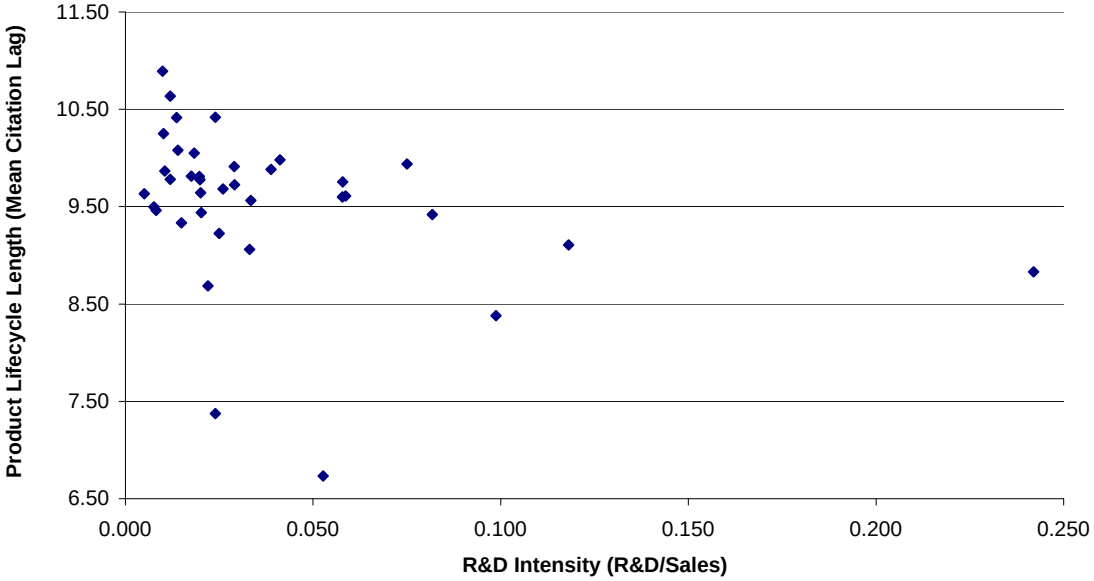

Notes: This figure plots the product lifecycle length against the R&D intensity of each SIC 3-digit industry studied in this paper. Product lifecycle lengths are average patent citation lags and are computed using data in the NBER Patent Citation Datafile. R&D intensity is the R&D-to-sales ratio in the BEA firm-level data during the benchmark years between 1982 and 2004.

Figure 6: Flexible Estimation - Number of Affiliates

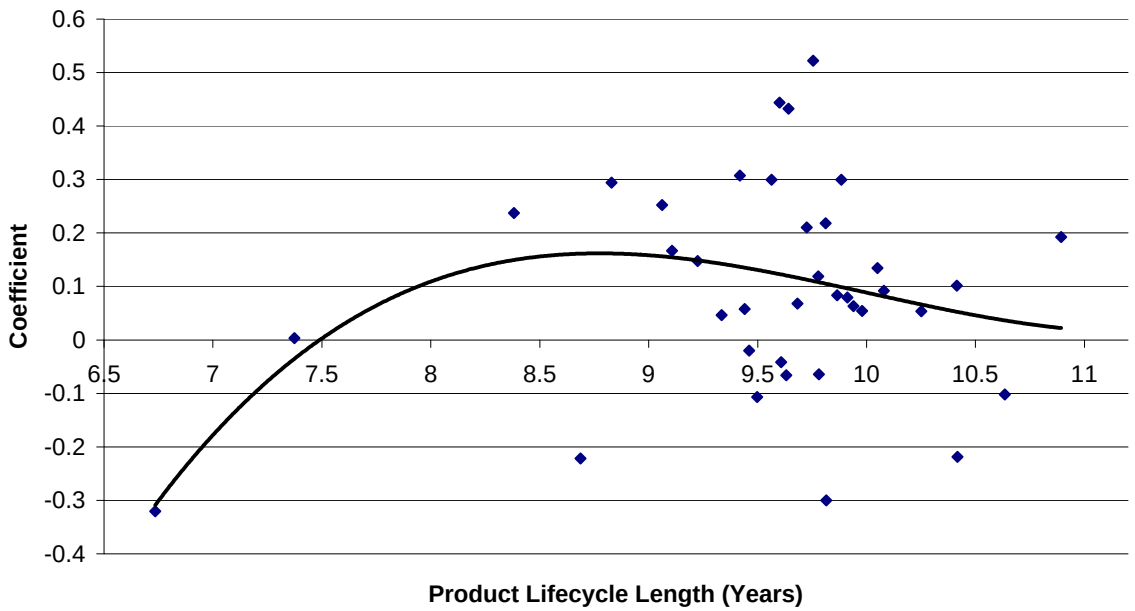

Notes: This figure plots estimated coefficients from an OLS regression similar to (21) against product lifecycle lengths, but with a separate interaction for each of the 37 industries. The dependent variable is the log number of affiliates.

Figure 7: Flexible Estimation - Affiliate Presence

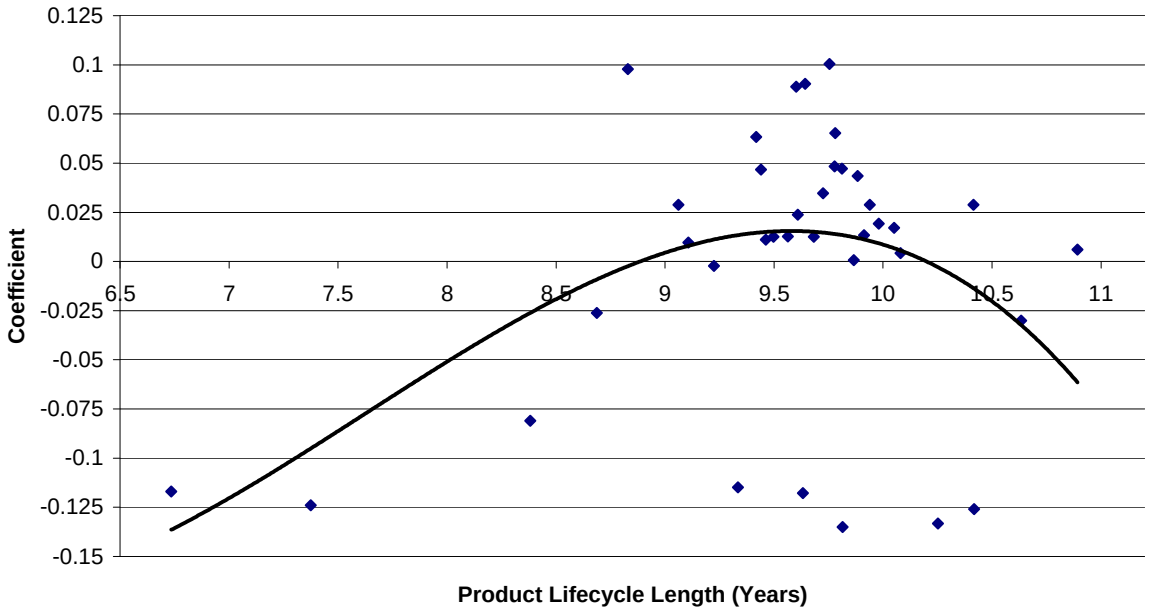

Notes: This figure plots estimated coefficients from an OLS regression similar to (21) against product lifecycle lengths, but with a separate interaction for each of the 37 industries. The dependent variable is an indicator for positive affiliate sales.

Table 1: Regression Summary Statistics

Variable	Mean	Standard Deviation	Min	Max
Country-Industry-Year Level:				
Total Affiliate Sales	257000	2000000	-9100	88600000
Total Affiliate Employment	942	5300	0	161000
Total Affiliate Assets	316000	3120000	0	171000000
Total Affiliate R&D	3300	42800	0	2131732
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:				
Average Affiliate Sales	80200	491000		
Average Affiliate Employment	310	1290		
Average Affiliate Assets	79500	439000		
Average Affiliate R&D	1070	16400		
Average Affiliate Sales to Local Unaffiliated (only majority-owned)	43000	210000		
Average Affiliate Sales to U.S. (only majority-owned)	10900	214000		
Industry Level:				
Average Patent Citation Lag (years)	9.55	0.79	6.73	10.9
Average 85 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
U.S. exports	74300	618000	0	39000000
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	1125
Number of Affiliates per Parent	9.74	20.5		
Number of Countries	92			
Number of Industries	37			
Number of Observations, Industry-level	15540			
Number of Observations, Industry-level, Conditional on Presence	4977			
Number of Observations, Affiliate-level	22505			

Notes: This table summarizes multinational activity, host-country institutions, and industry characteristics across 92 countries, 37 industries, and all benchmark years during 1982-2004. All financial values are reported in thousands of dollars. Average patent citation lags were calculated using the NBER U.S. Patent Citations Data File (Hall, et al 2001) by patent class and matched to 3-digit SIC industry codes using a standard USPTO concordance. GDP per capita is from the Penn World Table (Heston, et al 2009). U.S. exports are from the Census Bureau (http://www.som.yale.edu/faculty/pks4/). All other variables are from the Bureau of Economic Analysis Survey of U.S. Direct Investment Abroad, and pertain to U.S. outward foreign direct investment.

Table 2: Product Lifecycle Lengths and R&D Intensities by Sector, Ranked by Turnover

Industry Name	Ranking	SIC Code	Average Citation Lag (Years)	R&D Intensity
Shortest Citation Lags:				
Electronics Machinery	1	383	6.73	.0527
Watches, Clocks, Clockwork Operated Devices	2	387	7.37	.0239
Computer And Office Equipment	3	357	8.38	.0987
Agricultural Chemicals	4	287	8.69	.0219
Electronic Components And Accessories	5	367	8.83	.242
75th-Percentile Citation Lags:				
Paints, Varnishes, Lacquers, Enamels	25	285	9.81	.0180
Rolling, Drawing, Extruding Of Metals	26	335	9.87	.0105
Miscellaneous Electrical Machinery, Equipment	27	369	9.88	.0388
Engines And Turbines	28	351	9.91	.0290
Communications Equipment	29	366	9.94	.0750
Longest Citation Lags:				
Fabricated Structural Metal Products	33	344	10.25	.0102
Cutlery, Handtools, And General Hardware	34	342	10.41	.0137
Screw Machine Products, Bolts, Nuts, Screws	35	345	10.42	.0240
Metal Cans And Shipping Containers	36	341	10.63	.0119
Heating Equipment, Except Electric	37	343	10.89	.00986

Notes: This table shows the average patent citation lag and R&D intensity for the top and bottom five industries ranked by average citation lags. Citation lags were calculated using the NBER U.S. Patent Citations Data File (Hall, et al 2001). To preserve comparability across patents granted at different dates, citations were limited to within 15 years of the patent grant date. Patent classes were matched to 3-digit SIC industry codes a standard USPTO concordance, available from ftp://ftp.uspto.gov/pub/taf/sic_conc/2005_diskette/. 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.

Table 3: Host-country Patent Laws and Affiliate Presence, Industry Level

Dependent variable:	Indicator for positive sales						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	-0.0009	-0.2263	-2.0279	-0.1452	-0.7672	-0.2773	-0.4722
	0.0101	0.045***	0.2738***	0.0451***	0.3268**	0.05***	0.3183
IPR x T		0.0236	0.4341	0.0151	0.1569	0.0264	0.0709
		0.0048***	0.0616***	0.005***	0.0742**	0.0054***	0.0716
IPR x T^2			-0.0231		-0.008		-0.0025
			0.0034***		0.0042*		0.004
IPR x R&D Intensity						0.6388	0.8471
						0.0893***	0.255***
IPR x R&D Intensity^2							-1.5269
							0.9303
log GDP per Capita	0.0925	0.0925	0.0925	-0.0477	-2.0855	-0.0477	-2.0855
	0.0513*	0.0514*	0.0514*	0.0704	0.4015***	0.0704	0.4015***
log GDPpc x T				0.0147	0.479	0.0147	0.479
				0.0047***	0.0908***	0.0047***	0.0908***
log GDPpc x T^2					-0.0261		-0.0261
					0.0051***		0.0051***
Controls	Country, Year, and Industry FE, Corporate Tax Rate						
N	15281	15281	15281	15281	15281	15281	15281
R^2	0.5216	0.5235	0.5268	0.5239	0.5299	0.5276	0.5322

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (18). The dependent variable indicates positive sales by affiliates of U.S.-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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends, country-year fixed effects, and sector-year fixed effects The results shown above were estimated with OLS (Angrist and Pischke 2009), and nearly identical results obtain with probit estimation.

Table 4: Host-country Patent Laws and Affiliate Activity, Industry Level

Dependent variable:	Log affiliate sales			Log affiliate assets			Log affiliate employment		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
IPR	0.0898	-8.1725	-8.5096	0.0368	-9.2692	-10.047	0.0183	-8.0969	-8.2497
	0.0685	3.6502**	3.6902**	0.0595	3.9716**	4.0068**	0.0716	3.3539**	3.2998**
IPR x T		1.6857	1.6741		1.9406	1.9518		1.6307	1.6007
		0.7547**	0.764**		0.8209**	0.8248**		0.6962**	0.6865**
IPR x T^2		-0.0856	-0.0818		-0.1008	-0.0954		-0.0815	-0.0777
		0.0392**	0.0397**		0.0424**	0.0425**		0.0363**	0.0359**
IPR x R&D Intensity			3.1351			5.2001			3.1174
			1.1883***			1.2684***			1.218**
IPR x R&D Intensity^2			-9.966			-14.8333			-12.364
			4.8455**			4.9555***			5.3307**
log GDP per Capita	1.545	14.277	13.1172	1.6091	14.2245	12.0387	0.7717	19.6372	18.9507
	0.3887***	6.8514**	6.8379*	0.000***	7.6158*	7.5242	.4723*	7.7283**	7.6158**
log GDPpc x T		-2.3399	-2.1076		-2.3449	-1.9062		-3.6423	-3.5044
		1.4385	1.4352		1.5426	1.5249		1.6171**	1.5953**
log GDPpc x T^2		0.1048	0.093		0.1066	0.0842		0.1738	0.1669
		0.0741	0.074		0.0788	0.078		0.0831**	0.0821**
Controls	Country, Year, and Industry FE, Corporate Tax Rate								
N	4952	4952	4952	4964	4964	4964	4788	4788	4788
R^2	0.6264	0.63	0.6307	0.6498	0.6536	0.656	0.576	0.5797	0.5804

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (18). The dependent variable is the log of sales (columns 1-3), the log of assets (columns 4-6), or the log of employment (columns 7-9), of affiliates of U.S.-based multinational firms by country, sector, and year. These outcome variables are 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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends, country-year fixed effects, and sector-year fixed effects.

Table 5: Host-country Patent Laws and Affiliate Activity, Industry Level, First Differences

Dependent variable:	Indicator for increased sales						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Delta IPR	0.0629	-0.1622	-5.3506	-0.1131	-4.5372	-0.5129	-3.627
	0.0218***	0.0525***	0.4964***	0.0462**	0.5829***	0.0702***	0.5529***
Delta IPR x T		0.0236	1.2058	0.0184	1.0265	0.0526	0.7429
		0.0055***	0.1129***	0.0049***	0.1341***	0.0072***	0.1261***
Delta IPR x T^2			-0.0665		-0.0567		-0.0384
			0.0063***		0.0075***		0.0071***
Delta IPR x R&D Intensity						1.9329	4.3895
						0.2164***	0.4864***
Delta IPR x R&D Intensity^2							-12.925
							1.6376***
Delta log GDP per Capita	0.2848	0.2848	0.2848	-0.2461	-8.5075	-0.2461	-8.5075
	0.1126**	0.1126**	0.1126**	0.1164**	3.9134**	0.1164**	3.9137**
Delta log GDPpc x T				0.0556	1.9381	0.0556	1.9381
				0.0148***	0.9074**	0.0148***	0.9075**
Delta log GDPpc x T^2					-0.1058		-0.1058
					0.0506**		0.0506**
Controls	Year FE, Delta Corporate Tax Rate						
N	11803	11803	11803	11803	11803	11803	11803
R^2	0.0489	0.0497	0.0611	0.0502	0.0647	0.0642	0.0779

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (19). The dependent variable indicates an increase in sales by affiliates of U.S.-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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends and sector-year fixed effects. The results shown above were estimated with OLS (Angrist and Pischke 2009), and nearly identical results obtain with probit estimation.

Table 6: Host-country Patent Laws and Number of Affiliates, Industry Level

Dependent variable:	Log number of affiliates						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	0.0891	0.6238	-3.8085	-0.5244	-5.283	-0.922	-5.6723
	0.035**	0.1989***	1.2946***	0.3154	2.0246**	0.341***	1.9728***
IPR x T		-0.0557	0.8835	0.0645	1.0741	0.102	1.0883
		0.0204***	0.2759***	0.0328*	0.4163**	0.0348***	0.4067***
IPR x T^2			-0.0496		-0.0534		-0.0514
			0.0146***		0.0214**		0.021**
IPR x R&D Intensity						0.8429	2.2488
						0.2805***	0.5743***
IPR x R&D Intensity^2							-6.1811
							2.3459**
log GDP per Capita	0.8249	0.8224	0.8179	3.1424	4.9386	3.1005	4.0017
	0.2076***	0.207***	0.2066***	0.5677***	3.7801	0.5713***	3.7732
log GDPpc x T				-0.2451	-0.6283	-0.2422	-0.4403
				0.057***	0.774	0.0575***	0.7728
log GDPpc x T^2					0.0203		0.0107
					0.0399		0.0399
Controls	Country, Year, and Industry FE, Corporate Tax Rate						
N	4977	4977	4977	4977	4977	4977	4977
R^2	0.6463	0.6465	0.6559	0.6466	0.6624	0.7126	0.7137

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (18). The dependent variable is the log of the number of affiliates of U.S.-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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends and country-year fixed effects.

Table 7: Host-country Patent Laws and Affiliate Activity, Firm-level

Dependent variable:	Log affiliate sales						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	0.0873	0.5898	-4.0966	0.3491	-7.733	0.4584	-7.9673
	0.0609	0.2685**	3.1421	0.4056	4.9015	0.4026	4.847
IPR x T		-0.0537	0.9473	-0.0282	1.6988	-0.0386	1.7012
		0.029*	0.6545	0.0419	1.0187*	0.0416	1.0134*
IPR x T^2			-0.0533		-0.092		-0.0906
			0.034		0.0529*		0.0529*
IPR x R&D Intensity						-0.1816	2.9381
						0.4625	0.7437***
IPR x R&D Intensity^2							-13.8194
							4.2491***
log GDP per Capita	0.853	0.8429	0.8479	1.4653	11.119	1.4762	11.8591
	0.2448***	0.2457***	0.2448***	0.8055*	9.811	0.8058*	9.7964
log GDPpc x T				-0.0665	-2.122	-0.0673	-2.2745
				0.0872	2.0494	0.0872	2.0492
log GDPpc x T^2					0.1092		0.1172
					0.1056		0.1057
Controls	Country, Year, Industry FE, Corporate Tax Rate Log Parent Sales and R&D						
N Observations	22505	22505	22505	22505	22505	22505	22505
R^2	0.2004	0.2006	0.2007	0.2007	0.2008	0.2007	0.2017

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (20). The dependent variable is the log of sales by affiliates of U.S.-based multinational firms by individual affiliate 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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982--2004. All regressions include country, year, and industry fixed effects, the log of parent-company sales, the log of parent-company R&D expenditures, host-country GDP per capita, and a measure of host-country corporate tax rates from the BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends.

Table 8: Host-country Patent Laws and Affiliate Activity, Firm-level, Firm Heterogeneity

Dependent variable	Log affiliate sales								
	U.S. Parent Productivity			U.S. Parent Productivity			U.S. Parent Productivity		
	Low	High	All firms	Low	High	All firms	Low	High	All firms
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
IPR	-12.2114	10.9268	-5.1439	-14.1185	2.9447	-18.0981	-14.5604	2.6936	-18.4396
	3.3223***	6.6321	3.2734	4.7783***	8.0843	5.2991***	4.775***	8.0042	5.2776***
IPR x T	2.6485	-2.2068	1.1228	3.1032	-0.586	3.9265	3.1461	-0.5747	3.9539
	0.7069***	1.3807	0.6961	0.9957***	1.6852	1.1033***	1.0038***	1.6739	1.1044***
IPR x T^2	-0.1421	0.1116	-0.0602	-0.1688	0.0297	-0.2114	-0.1694	0.0306	-0.2114
	0.0375***	0.0713	0.0369	0.0519***	0.0874	0.0575***	0.0528***	0.087	0.0579***
IPR x High Productivity			5.8788			24.699			24.6332
			1.4815***			8.3335***			8.3412***
IPR x T x High Productivity			-1.1457			-5.2651			-5.2539
			0.3164***			1.7367***			1.7373***
IPR x T^2 x High Productivity			0.0553			0.2799			0.2796
			0.0174***			0.0902***			0.0902***
IPR x R&D Intensity							2.9983	2.0103	2.7113
							0.7859***	1.1801*	0.7967***
IPR x R&D Intensity^2							-14.2084	-8.4809	-12.6312
							4.226***	5.6094	4.0268***
log GDP per Capita	0.975	0.7945	0.9109	6.5018	26.3546	20.9822	7.2094	26.7928	21.4414
	0.2598***	0.2906***	0.2398***	10.7325	15.1479*	11.3184*	10.7624	15.2519*	11.4003*
log GDPpc x T				-1.3095	-5.2083	-4.3786	-1.4448	-5.3039	-4.4737
				2.2702	3.1246	2.4122*	2.2791	3.1471*	2.4324*
log GDPpc x T^2				0.0766	0.2639	0.2378	0.0834	0.2689	0.2427
				0.1192	0.1593	0.1271*	0.1198	0.1605*	0.1284*
Controls	Country, Year, Industry FE, Corporate Tax Rate Log Parent Sales and R&D								
N Observations	9286	11052	20338	9286	11052	20338	9286	11052	20338
R^2	0.2083	0.2648	0.2386	0.2086	0.2657	0.2394	0.2097	0.266	0.2402

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (20) and its triple-differenced analog. Affiliates are categorized as subsidiaries of Low and High Productivity parent companies using a simple above-median Solow residual criterion. The dependent variable is the log of individual affiliate sales. T is the industry-specific average patent citation lag, R&D Intensity is the average R&D-to-sales ratio by industry, IPR is the index of patent strength, and GDP per Capita (GDPpc) is per-capita GDP. I include but do not report interactions between High Productivity and variables involving GDP per capita and R&D intensity. The benchmark survey data are from the BEA and were collected in 1982, 1989, 1994, 1999, and 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends.

Table 9: Exports Versus Multinational Activity

Dependent variable:	Fraction of sales by multinational affiliates						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR	-0.0002	-0.1142	-2.1914	-0.1364	-1.402	-0.2486	-1.132
	0.0092	0.0474**	0.2966***	0.0511***	0.3696***	0.0533***	0.367***
IPR x T		0.0118	0.4629	0.0141	0.2891	0.0242	0.2163
		0.0049**	0.0668***	0.0054**	0.0835***	0.0056***	0.0832**
IPR x T^2			-0.0243		-0.0148		-0.0104
			0.0037***		0.0047***		0.0047**
IPR x R&D Intensity						0.3894	0.592
						0.064***	0.2235***
IPR x R&D Intensity^2							-1.2662
							0.8538
log GDP per Capita	0.0281	0.0284	0.0283	0.0681	-1.4166	0.065	-1.4194
	0.0195	0.0193	0.0194	0.0686	0.5148***	0.068	0.5152***
log GDPpc x T				-0.0041	0.3174	-0.0038	0.3178
				0.0066	0.1145***	0.0066	0.1146***
log GDPpc x T^2					-0.0173		-0.0173
					0.0064***		0.0064***
Controls	Country, Year, and Industry FE, Corporate Tax Rate						
N	12651	12651	12651	12651	12651	12651	12651
R^2	0.5237	0.5244	0.5274	0.5245	0.5281	0.5267	0.5297

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (18). The dependent variable is the ratio of affiliate sales to the sum of affiliate sales plus U.S. exports by country, sector, and year, for affiliates of U.S.-based multinationals, and is based on firm-level data from the BEA and export data from the U.S. Census Bureau. IPR is the index of patent protection from Ginarte and Park (1997) and Park (2008). T is the product lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends, country-year fixed effects, and sector-year fixed effects.

Table 10: Flexible Estimation

Dependent variable:	Indicator for any affiliates				Log number of affiliates			
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IPR	-0.0151	-0.0208	-0.0435	-0.0493	0.1562	0.042	-0.0034	-0.1061
	0.0127	0.0109*	0.0145***	0.0126***	0.0574***	0.0432	0.0524	0.0451**
IPR x T_2	0.0591	0.0416	0.0618	0.0442	-0.0111	0.0677	0.0195	0.0964
	0.0099***	0.0113***	0.0112***	0.0124***	0.0293	0.0409	0.0249	0.0398**
IPR x T_3	0.0315	0.033	0.04	0.0416	0.0438	0.1177	0.0895	0.1589
	0.006***	0.0078***	0.0061***	0.0077***	0.0213**	0.0348***	0.0234***	0.0376***
IPR x T_4	0.0007	0.0105	0.0147	0.0245	-0.1072	0.0483	-0.0191	0.1302
	0.0093	0.0105	0.0086*	0.0103**	0.0277***	0.0365	0.0231	0.0367***
log GDPpc x T Quartiles	N	Y	N	Y	N	Y	N	Y
IPR x R&D Intensity Quartiles	N	N	Y	Y	N	N	Y	Y
Corporate Tax Rate and GDPpc	Y	Y	Y	Y	Y	Y	Y	Y
Country, Year, and Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
N	15540	15540	15540	15540	4916	4916	4916	4916
R^2	0.5295	0.5313	0.5309	0.5328	0.7039	0.7142	0.7088	0.7184

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (21). The dependent variable is based on the number of affiliates of U.S.-based multinational firms by country, sector, and year based on firm-level data from the BEA, and is an indicator for any affiliates (columns 1-4) or the log of the number of affiliates (columns 5-8). IPR is the index of patent protection from Ginarte and Park (1997) and Park (2008). T_2, T_3, and T_4 are the 2nd, 3rd, and 4th quartiles of the product lifecycle length distribution, by industry, based on data from the USPTO. 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 et al (2009). The sample period is 1982--2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and a measure of host-country corporate tax rates based on BEA data. Standard errors, adjusted for clustering at the country-level, appear below each point estimate. Standard errors, adjusted for clustering at the country-level, appear below each point estimate.

Table 11: Patent Effectiveness, Anticipated Reforms, and Host-Country Credit Conditions

Panel A: Patent Effectiveness

Dependent variable:	Log affiliate sales				Log number of affiliates			
	Patent Effectiveness		Patent Effectiveness		Patent Effectiveness		Patent Effectiveness	
	Low	High	Low	High	Low	High	Low	High
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IPR	0.055	-0.8321	4.3335	-10.2991	0.9037	-1.1308	10.0963	-14.2651
	0.4871	0.3199**	4.6496	2.21***	0.8894	0.4563**	8.8397	3.7315***
IPR x T	0.0016	0.0982	-0.8879	2.1329	-0.087	0.1289	-1.9964	2.9456
	0.0499	0.0344***	0.9469	0.4577***	0.0921	0.0485***	1.8063	0.7714***
IPR x T^2			0.0461	-0.1088			0.0989	-0.1503
			0.0482	0.0236***			0.0924	0.0398***
Controls	Country, Year, and Industry FE, Corporate Tax Rate, log GDP per Capita and Interactions							
N	2281	2696	2281	2696	2269	2683	2269	2683
R^2	0.7087	0.7277	0.7104	0.7328	0.6157	0.6612	0.6166	0.6624

Panel B: Anticipated Reforms

Dependent variable:	Indicator for positive sales			Log affiliate sales			Log number of affiliates		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
IPR	-0.1892	-0.6489	-0.3379	-0.6394	-4.3688	-4.5616	-0.873	-7.647	-7.7959
	0.056***	0.3568*	0.3489	0.3143**	2.2376*	2.1772**	0.4875*	3.8219**	3.8668**
IPR x T	0.0203	0.125	0.0342	0.0726	0.8654	0.8394	0.0935	1.5322	1.4764
	0.0058***	0.0795	0.0769	0.0326**	0.4671*	0.4569*	0.0499*	0.7997*	0.8064*
IPR x T^2		-0.0059	-0.0001		-0.042	-0.0381		-0.0761	-0.0699
		0.0044	0.0042		0.0244*	0.024		0.0419*	0.0422
R&D Intensity Interactions	N	N	Y	N	N	Y	N	N	Y
Controls	Country, Year, and Industry FE, Corporate Tax Rate, log GDP per Capita and Interactions								
N	11877	11877	11877	4287	4287	4287	4265	4265	4265
R^2	0.541	0.5475	0.5504	0.7294	0.73	0.7317	0.6583	0.6585	0.6594

Panel C: Placebo Test: Host-Country Credit Conditions

Dependent variable:	Indicator for positive sales			Log affiliate sales			Log number of affiliates		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Credit	-0.0141	-0.2836	-0.1495	0.0145	2.8669	2.7396	-0.1889	2.4037	1.7705
	0.0505	0.2442	0.2379	0.2519	1.8073	1.8356	0.3874	3.0938	3.165
Credit x T	-0.0004	0.061	0.0199	-0.0026	-0.6083	-0.611	0.0279	-0.5231	-0.4936
	0.0051	0.0557	0.0541	0.0238	0.3741	0.3817	0.0384	0.6396	0.6703
Credit x T^2		-0.0035	-0.0008		0.0321	0.0334		0.0292	0.0318
		0.0032	0.0031		0.0193	0.0198*		0.0329	0.0353
R&D Intensity Interactions	N	N	Y	N	N	Y	N	N	Y
Controls	Country, Year, and Industry FE, Corporate Tax Rate, log GDP per Capita and Interactions								
N	13579	13579	13579	4908	4908	4908	4889	4889	4889
R^2	0.5259	0.5325	0.5332	0.7119	0.7124	0.7127	0.635	0.6351	0.6364

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports three extensions of the main results. The top panel categorizes industries by patent effectiveness, a measure obtained from Cohen, et al (2000); patent effectiveness is considered High for sectors with above-median scores. The middle panel shows estimation results comparable to those shown in Tables 3 and 4, but observations known to have involved pre-announced reforms associated with TRIPS (Kyle and McGahon 2009) have been omitted from the sample. The lower panel shows the results of a placebo test in which a time-varying measure of host-country credit conditions is used in place of the patent index. The index is from Djankov, et al (2007). In all regressions, T is the industry-specific product lifecycle length and IPR is the index of patent protection. The sample period is 1982-2004. Standard errors, adjusted for clustering at the country-level, appear below each point estimate.

Table 12: Growth in Affiliate Activity, Industry Level

Dependent variable:	log(Sales$_{t+1}$) - log(Sales$_t$)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IPR$_t$	-0.0187	-1.1595	-14.6835	-1.2995	-16.3583	-1.0157	-15.8938
	0.0699	0.2628***	3.2119***	0.4336***	5.3521***	0.5128*	5.1661***
IPR$_t$ x T		0.1188	2.9614	0.1334	3.2915	0.1064	3.1924
		0.0273***	0.6747***	0.0466***	1.1195***	0.0536*	1.0835***
IPR$_t$ x T^2			-0.1489		-0.1651		-0.16
			0.0355***		0.0585***		0.0568***
IPR$_t$ x R&D Intensity						-0.5326	0.7051
						0.4225	1.0574
IPR$_t$ x R&D Intensity^2							-4.1074
							4.774
log GDP per Capita$_t$	-1.0658	-1.0623	-1.0818	-0.7603	2.2582	-0.7405	2.0862
	0.3139***	0.3127***	0.3078***	0.7232	6.9763	0.725	6.9627
log GDPpc$_t$ x T				-0.0314	-0.655	-0.0326	-0.6198
				0.0713	1.4133	0.0713	1.4095
log GDPpc$_t$ x T^2					0.032		0.0302
					0.0726		0.0724
Controls	Country, Year, and Industry FE, Corporate Tax Rate$_t$						
N	3280	3280	3280	3280	3280	3280	3280
R^2	0.1253	0.1278	0.1305	0.1278	0.1306	0.1281	0.1308

Notes: * p<0.10, ** p<0.05, *** p<0.01. This table reports least-squares estimates of equation (18). The dependent variable is the forward growth in the log of affiliate sales over one period for affiliates of U.S.-based multinational firms by country, sector, and year 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 lifecycle length, by industry, and is the average patent citation lag based on data from the USPTO. 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 et al (2009). The sample period is 1982-2004. All regressions include country, year, and industry fixed effects, host-country GDP per capita, and host-country corporate tax rates based on BEA data. 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 U.S. outward FDI, China and India, and the chemical and pharmaceutical industries, as well as including country-specific time trends and country-year fixed effects.

1 Appendix

A.1 Theory

This section includes detailed derivations of theoretical results appearing in the main body of the paper.

Imitation incentives, equations (6)-(7): Southern entrepreneurs will pursue imitation of a sector-j variety of maturity t if the net present value $V_0(\xi_S, t, j)$ of doing so exceeds cost c,

$$V_0(\xi_S, t, j) = E_m \left[\int_0^{\max\{0, T_j - t - m\}} e^{-rs} v(s) ds \right].$$

Assuming there is no time discounting, and noting that the flow value $v(s) = (1 + \xi_S)\pi_{im}^S$ is independent of time, we can obtain the following,

$$\begin{aligned} V_0(\xi_S, t, j) &= E_m \left[\int_0^{\max\{0, T_j - t - m\}} (1 + \xi_S)\pi_{im}^S ds \right] \\ &= (1 + \xi_S)\pi_{im}^S E_m \left[\max\{0, T_j - t - m\} \right]. \end{aligned}$$

Computing the expected value at right,

$$\begin{aligned} E_m[\max\{0, T_j - t - m\}] &= E_m[(T_j - t - m)1\{m < T_j - t\}] + E_m[0 \cdot 1\{m \geq T_j - t\}] \\ &= \int_0^{T_j - t} (T_j - t - m) f(m) dm \\ &= \frac{(T_j - t)^2}{\overline{m}} - \int_0^{T_j - t} \frac{m}{\overline{m}} dm \\ &= \frac{(T_j - t)^2}{2\overline{m}}, \end{aligned}$$

we arrive at the net present value of imitation, equation (6)

$$V_0(\xi_S, t, j) = (1 + \xi_S)\pi_{im}^S \cdot \frac{(T_j - t)^2}{2\overline{m}}.$$

Sourcing incentives, equations (8)-(11): Given ξ_S, a sector-j firm will begin offshoring manufacturing activity at product maturity t to maximize expected lifetime profit,

$$\begin{aligned} E_m[\Pi_j(t)] &= \int_0^t 2\pi^N ds + E_m \left[\int_t^{T_j} [2\pi^S \cdot 1\{t + m > s\} + (1 + \xi_S)\pi^S \cdot 1\{t + m \leq s\}] ds \right] \\ &= \begin{cases} 2\pi^N t + 2\pi^S E[\min\{T_j - t, m\}] \\ \qquad\qquad + (1 + \xi_S)\pi^S E[\max\{0, T_j - t - m\}], & \text{if } t \leq T_j - \overline{\tau}(\xi_S) \\ 2\pi^N t + 2\pi^S (T_j - t), & \text{if } t > T_j - \overline{\tau}(\xi_S). \end{cases} \end{aligned}$$

Notice that the expectations in the first case are

$$\begin{aligned}
E_m[\max\{0, T_j - t - m\}] &= \frac{(T_j - t)^2}{2\overline{m}} \text{ as above, and} \\
E_m[\min\{m, T_j - t\}] &= (T_j - t) \cdot P\{m \geq T_j - t\} + E[m \cdot 1\{m < T_j - t\}] \\
&= (T_j - t) \cdot \left(1 - \frac{T_j - t}{\overline{m}}\right) + \int_0^{T_j - t} m \cdot f(m) dm \\
&= T_j - t - \frac{(T_j - t)^2}{2\overline{m}}.
\end{aligned}$$

The firm's objective in the first case is therefore

$$E_m[\Pi_j(t)] = 2\pi^N t + 2\pi^S \left(T_j - t - \frac{(T_j - t)^2}{2\overline{m}}\right) + (1 + \xi_S)\pi^S \left(\frac{(T_j - t)^2}{2\overline{m}}\right).$$

Optimizing over the initial offshoring maturity t, it is apparent that

$$\begin{aligned}
0 &= 2\pi^N + 2\pi^S \left(\frac{T_j - t}{\overline{m}} - 1\right) - (1 + \xi_S)\pi^S \frac{T_j - t}{\overline{m}} \\
\Rightarrow 2(\pi^S - \pi^N) &= \frac{T_j - t}{\overline{m}}(1 - \xi_S)\pi^S \\
\Rightarrow \tau^*(\xi_S) \equiv T_j - t^* &= \frac{\pi^S - \pi^N}{(1 - \xi_S)\pi^S} 2\overline{m}.
\end{aligned}$$

The parameter values of the model dictate whether the firm does better by following this strategy (i.e. offshoring at $\tau^*(\xi_S)$ as defined above), which subjects the firm to imitation risk, or whether it is preferable to take the relatively conservative strategy of case 2, offshoring at $\overline{\tau}(\xi_S)$ and avoiding imitation risk. Comparing lifetime profit under the two strategies, it is apparent as described in equation (9) that bearing some imitation risk (case 1) is preferred, provided that the cost benefits of manufacturing in the South are high relative to the impact of imitation. In either case, firms follow a sector-invariant sourcing rule, whereby production is offshored at a time-to-obsolescence that is a function of Southern intellectual property protection ξ_S. This result implies a measure of manufacturing activity in the North $N_j(\xi_S)$ that is increasing in T_j at any ξ_S: $\frac{\partial N_j(\xi_S)}{\partial T_j} \geq 0$, which is Proposition 1.

The distribution of revenues earned by Southern affiliates across sectors j is also determined by the sourcing cut-off described by equation (9),

$$R_j(\xi_S) = \int_{\max\{0, T_j - \tau^*(\xi_S)\}}^{T_j} \left(2r^S(1 - \kappa_{im}(t)) + (1 + \xi_S)r^S \kappa_{im}(t)\right) \psi_j(t) dt, \tag{1}$$

where $\psi_j(t) = 1/T_j$ is the density of product maturities and $\kappa_{im}(t)$ is the probability that a maturity-t product has been imitated. It is straightforward to see that if product lifecycles are so short that firms face no imitation risk (i.e. $T_j < \overline{\tau}(\xi_S)$), Southern affiliates will perform all manufacturing activity and earn $2r^S$. For firms in sectors that face imitation risk but nevertheless manufacture in the South at all maturities (i.e. $T_j \in [\overline{\tau}(\xi_S), \tau^*(\xi_S)]$), some measure of varieties

will be imitated at any point in time: $\kappa_{im}(t) = \frac{t}{\overline{m}}$. This implies that revenues earned by Southern affiliates will be

$$\begin{aligned} R_j(\xi_S) &= \int_0^{T_j} \left(2r^S(1-\kappa_{im}(t)) + (1+\xi_S)\cdot r^S\kappa_{im}(t)\right)\psi_j(t)dt \\ &= \int_0^{T_j} \left(2r^S(1-\frac{t}{\overline{m}}) + (1+\xi_S)\cdot r^S\frac{t}{\overline{m}}\right)\frac{1}{T_j}dt \\ &= 2r^S\left(1-\frac{T_j}{2\overline{m}}\right) + (1+\xi_S)r^S\frac{T_j}{2\overline{m}}. \end{aligned}$$

For firms in longer-lifecycle sectors (with $T_j > \tau^*(\xi_S)$), only the fraction with relatively mature products will manufacture in the South. Of this fraction, a subset of varieties will be imitated at any point in time: $\kappa_{im}(t) = \frac{t-[T_j-\tau^*(\xi_S)]}{\overline{m}}$. Total Southern-affiliate revenues in this case are

$$\begin{aligned} R_j(\xi_S) &= \int_{T_j-\tau^*(\xi_S)}^{T_j} \left(2r^S(1-\kappa_{im}(t)) + (1+\xi_S)\cdot r^S\kappa_{im}(t)\right)\psi_j(t)dt \\ &= \int_{T_j-\tau^*(\xi_S)}^{T_j} \left[2r^S\left(1-\frac{t-[T_j-\tau^*(\xi_S)]}{\overline{m}}\right) + (1+\xi_S)\cdot r^S\left(\frac{t-[T_j-\tau^*(\xi_S)]}{\overline{m}}\right)\right]\psi_j(t)dt, \end{aligned}$$

which, after a change of variables $\widetilde{t} = t - [T_j - \tau^*(\xi_S)]$ can be reduced to

$$\begin{aligned} R_j(\xi_S) &= \int_0^{\tau^*(\xi_S)} \left[2r^S(1-\frac{\widetilde{t}}{\overline{m}}) + (1+\xi_S)\cdot r^S\frac{\widetilde{t}}{\overline{m}}\right]\frac{1}{T_j}d\widetilde{t} \\ &= \frac{1}{T_j}\left[(2r^S\left(\tau^*(\xi_S) - \frac{\tau^*(\xi_S)^2}{2\overline{m}}\right) + (1+\xi_S)r^S\frac{\tau^*(\xi_S)^2}{2\overline{m}}\right]. \end{aligned}$$

We thus arrive at the steady-state distribution of affiliate revenues across sectors, equation (11).

Response to reforms, equations (12)-(13): When the South strengthens intellectual property rights from ξ_S to ξ'_S, the sourcing cut-off rises: $\tau^*(\xi_S) \nearrow \tau^*(\xi'_S)$. As a result, firms offshore to the South manufacturing activity associated with *marginal* varieties. The remaining time to obsolescence τ of a marginal variety is such that it would not yet be offshored under the old patent regime, but will be under the strengthened regime: $\tau^*(\xi_S) < \tau < \tau^*(\xi'_S)$. The measure of varieties manufactured in the North is thus higher under the old patent regime, as described by equation (12) and as plotted in the middle panel of Figure 3.

Similarly, revenues earned by Southern affiliates will rise following the patent reform. This reflects the increased measure of varieties manufactured there as well as an improvement in the patent protection provided to existing imitated varieties. These distinct responses combine to generate an overall increase in observed affiliate revenues as described by equation (13). The pattern of response across sectors is an implication of the steady-state distribution of affiliate revenues described by (11). Notice that $\overline{\tau}(\xi_S) < \tau^*(\xi_S) < \tau^{*\prime}(\xi_S)$, so that each line of (13) is the difference $R_j(\xi'_S) - R_j(\xi_S)$. It is simple to show that analogous predictions hold in the cross section of countries with different levels of intellectual property rights protection.

Sector-level affiliate presence (0-1 margin), section 5.1: In a three-world country with transport costs and plant-level economics of scale, a Southern patent reform may attract multinational affiliates to a country-sector that previously had no foreign direct investment. As will be shown below, this effect is more pronounced in sectors with long product lifecycles.

Consider three countries: N, $\underline{S}$, and $\overline{S}$. The North N has wages w^N and patent protection $\xi_N = 1$ while the two Southern countries $\underline{S}$ and $\overline{S}$ share identical wages w^S and patent institutions $\xi^S < 1$; as before, $w^N > w^S$. Following the approach in Helpman, Melitz, and Yeaple (2004), I assume that transport costs are of the iceberg type and that firms face a fixed cost of establishing a foreign affiliate; I also assume that this fixed cost differs across countries. Specifically, suppose the cost of establishing an affiliate is $\underline{f}_I$ in country $\underline{S}$, and is $\overline{f}_I$ in country $\overline{S}$, that $\overline{f}_I > \underline{f}_I$, and that an exporting firm or affiliate must ship $\gamma > 1$ units for one unit to arrive at its destination.

Imitation incentives, as before, lead to a time-to-obsolescence cut-off $\overline{\tau}(\xi_S)$, similar to that described in (7) and increasing in Southern patent protection ξ_S as before. Northern innovators, however, now determine both the product maturity at which to begin offshoring production, and how to organize offshore activity. In particular, firms must decide between 1) establishing one offshore affiliate that serves both Southern markets, and 2) establishing two offshore affiliates, one in each Southern country. Which of these two strategies is optimal depends on the firm's objective function, similar to equation (8).

Specifically, given ξ_S, a sector-j firm will begin offshoring manufacturing activity at product maturity t to maximize expected lifetime profit,

$$E_m[\Pi_j(t)] = \int_0^t \pi_0 ds + \max\left\{E_m\left[\int_t^{T_j} A_{1j}(s)ds\right], E_m\left[\int_t^{T_j} A_{2j}(s)ds\right]\right\},$$

where net post-t profit with one offshore affiliate $A_{1j}(s)$ or two offshore affiliates $A_{2j}(s)$ is

$$\begin{aligned} A_{1j}(s) &= 1\{t+m>s\}\cdot\pi_1 + 1\{t+m\le s\}\cdot\pi_1(\xi_S) - \underline{f} \\ A_{2j}(s) &= 1\{t+m>s\}\cdot\pi_2 + 1\{t+m\le s\}\cdot\pi_2(\xi_S) - \underline{f} - \overline{f}, \end{aligned}$$

and m is the random time required for imitation. In the expressions above, gross profit earned by an unimitated firm with zero, one, or two offshore affiliates (π_0, π_1, and π_2, respectively) is

$$\begin{aligned} \pi_0 &= \lambda\left(\frac{\rho}{w^N}\right)^{\sigma-1}(1-\rho)(1+2\gamma^{1-\sigma}) \\ \pi_1 &= \lambda\left(\frac{\rho}{(w^N)^\alpha(w^S)^{1-\alpha}}\right)^{\sigma-1}(1-\rho)(1+2\gamma^{1-\sigma}) \\ \pi_2 &= \lambda\left(\frac{\rho}{(w^N)^\alpha(w^S)^{1-\alpha}}\right)^{\sigma-1}(1-\rho)(2+\gamma^{1-\sigma}), \end{aligned}$$

while gross profit earned by an imitated firm with one or two offshore affiliates ($\pi_1(\xi_S)$ and $\pi_2(\xi_S)$, respectively) is

$$\begin{aligned} \pi_1(\xi_S) &= \lambda\left(\frac{\rho}{w^N}\right)^{\sigma-1}(1-\rho)(\xi_S+(1+\xi_S)\gamma^{1-\sigma}) \\ \pi_2(\xi_S) &= \lambda\left(\frac{\rho}{(w^N)^\alpha(w^S)^{1-\alpha}}\right)^{\sigma-1}(1-\rho)(2\xi_S+\gamma^{1-\sigma}). \end{aligned}$$

For simplicity, suppose model parameters are such that firms follow the optimal offshoring strategy analogous to the first case in (9), initiating production in the South once the product is within $\tau^*(\xi_S) = \overline{\tau}(\xi_S)$ time to obsolescence. Then,

$$E_m[\Pi_j(t)] = \pi_0(T_j - \overline{\tau}(\xi_S)) + \max\left\{\overline{\tau}(\xi_S)\cdot\pi_1 - \underline{f}, \overline{\tau}(\xi_S)\cdot\pi_2 - \underline{f} - \overline{f}\right\}.$$

The firm will establish two Southern affiliates when offshoring first begins only if

$$\overline{\tau}(\xi_S)\cdot\pi_2 - \underline{f} - \overline{f} > \overline{\tau}(\xi_S)\cdot\pi_1 - \underline{f},$$

which is true only if the product has a sufficiently long remaining economic lifetime when offshoring first begins:

$$\overline{\tau}(\xi_S) > \overline{f}(\pi_2 - \pi_1)^{-1} \equiv F(w^N, w^S, \overline{f}).$$

For firms in sectors with $T_j > \overline{\tau}(\xi_S)$, an improvement in Southern patent protection from ξ_S to $\xi'_S > \xi_S$ leads to affiliate activity in $\overline{S}$ for the first time whenever $\overline{\tau}(\xi'_S) > F(w^N, w^S, \overline{f}) > \overline{\tau}(\xi_S)$. In such cases, however, firms in sectors with $T_j \leq \overline{\tau}(\xi_S)$ are unaffected, because products in these sectors do not have sufficiently long economic lifetimes to offset the fixed cost of establishing an affiliate in $\overline{S}$. Thus, an improvement in patent protection in the South may generate direct investment by multinationals in countries and sectors that previously had none, but this effect arises only in sectors with sufficiently long product lifecycles.

A.2 Data sources and variable construction

Product lifecycle lengths: In this paper, I construct a new measure of product lifecycle lengths based on information contained in the NBER Patent Citations Data File (Hall, et al 2001). A list of forward-citations was compiled for each U.S. patent granted between 1976 and 2006. For each patent, I compute the distribution of average within-15-year forward citation lags, and thereby obtain class-level moments including the mean, 75th percentile, and 85th percentile. My main proxy for product lifecycle lengths $\widehat{T}_j$ is the sector-specific mean. At the patent class level, $\widehat{T}_j$ has an average of 8.43 years and a standard deviation of 0.32. There is also variation across individual patents within each patent class, averaging 4.5 with a standard deviation of 0.45; the correlation between $\widehat{T}_j$ and this within-class variation is low, however, at 12.8%. Correlations between $\widehat{T}_j$ and other key sector-level characteristics at the SIC 3-digit level such as labor and capital intensity and measures of patent and secrecy effectiveness are described in the text, and are generally under 10% with the exception of R&D intensity, with which the correlation is -32%.

I also compare $\widehat{T}_j$ with alternative product lifecycle measures as described in section 3. In addition to 75th- and 85th-percentile citation lags and existing broadly-categorized measures of patent decay rates, I construct separate class-level proxies based on U.S. patent renewal data. U.S. utility patents issued on or after December 12, 1980 are subject to maintenance fees, which must be paid to keep the patent in force. Maintenance fees are due 3 1/2, 7 1/2 and 11 1/2 years from the date of the original patent grant (see http://www.uspto.gov/patents/process/maintain.jsp). Using USPTO maintenance fee data (available through Google; see http://www.google.com/googlebooks/uspto-patents-maintenance-fees.html), which records the dates of all maintenance fee events by patent, I compute each patent's duration, and thereby construct a distribution of patent durations by patent class. I then compute moments of this distribution including the mean, 75th percentile,

and 85th percentile durations. I find that the average renewal-based proxy is minimally correlated with $\widehat{T}_j$, while higher levels of correlation (approximately 30%) obtain for the for the 75th- and 85th-percentile proxies. Robustness checks based on the 85th-percentile renewal-based proxy also reveal that firms in high-T sectors are significantly more sensitive to patent protection than firms in fast-turnover sectors. Note, however, that there are important differences between my preferred measure and renewal-based measures: a key distinction is that $\widehat{T}_j$ is informed by technological links between patents, while renewal-based measures are limited to information contained only within each individual patent. Because firms' renewal strategies for individual patents on average differ across sectors for reasons unrelated to technology durability (for example, innovation success rates or the competitive conditions within an industry), this alternative proxy, much like the turnover rates in Broda and Weinstein (2010), may capture variation beyond that which T is meant to reflect in the model.

Patent rights index: The index of patent protection is published in Ginarte and Park (1997) and Park (2008). The index is available for 122 countries between 1960 and 2005, at five-year intervals, and is the sum of five sub-indexes corresponding to 1) enforcement, 2) coverage, 3) provisions for the loss of protection, 4) duration, and 5) membership in international intellectual property treaties. Further details are described at length in the two aforementioned publications. Based on additional empirical analysis, it is apparent that the most important components of the index for my results are the enforcement and membership sub-indexes. I thank Walter Park for generously providing me with the complete panel of sub-indexes.

Multinational activity: Confidential data on the activity abroad of U.S. multinational firms is provided by the Bureau of Economic Analysis through a sworn-status research arrangement. The data include detailed financial and operating information for each foreign affiliate owned (at least a 10% share) by a U.S. entity. The data variables used for this project were extracted from the BEA's comprehensive data files for each benchmark year, and then merged by parent and affiliate identification numbers to form a complete panel. Observations were excluded if a) values were carried over or imputed based on previous survey responses; b) the firm in question was in a sector that did not correspond to any patent class, according to the USPTO concordance described above; or c) the observation was a new entrant in the final year (2004) with a NAICS classification that could not be definitively matched to a SIC code in the industry sample. Of the approximately 55 sectors in the overall benchmark dataset, 37 primarily manufacturing sectors are included in my dataset; this corresponds to approximately 1000 U.S.-based parent companies per year, each with an average of ten foreign affiliate operations.