The Social Return of Research and Development Spending: Flux in Technology and Product Space Over Time

Carlos Carpi, Ian Lundberg and Rohan Thavarajah

May 5th 2014

Abstract

To what extent is a firm's performance influenced by the R&D activity of its neighbors? As Bloom et al (2013) explains, the key empirical challenge of this question is in separately quantifying the positive effect of knowledge diffusion from the negative effect of market rivalry. Bloom et al. (2013) address this challenge by assessing firm closeness along two separate dimensions: in product space (using sales activity broken down by industry) and technology space (using patenting activity across technology fields). We extend these findings by generating a time-varying measure of firm distances in technology space. We conclude that the strength of the positive effect of technology spillover derived in Bloom et al. (2013) hinges on the assumption that firm distances are static. In addition, we identify a source of asymmetry in the measurement error of distances in technology and product space. We find that after rescaling the market rivalry effect with our estimate of measurement error, the market rivalry effect overpowers that of technology diffusion so that firms must be closer in technology space than they are in product space for their R&D spending to net a positive effect on their peers.

Introduction

A firm's research and development investment position has an impact on its neighbors. As the firm contributes to a field's stock of knowledge, firms in the same field may utilize or build upon that contribution; we call this diffusion of knowledge or technology spillover. Although the R&D investment comes from individual firms, wider society reaps the benefits of technology spillovers. Policymakers might be interested in incentivizing R&D investment if R&D leads to high social returns beyond those experienced by the company. Quantifying spillover is of particular concern to policymakers, who would need to calibrate their subsidy policy to the size of this social return.

The effect of one firm's R&D expenditure on the performance of another is the net result of their interaction as technology neighbors in tech space and as rivals in product markets. This is the core empirical challenge of studies trying to identify spillover effects. As a firm increases its investment in R&D, we are interested in the uptick in the performance of its neighbors through tech spillovers. However, as a firm increases its investment in R&D, that firm also tends to become more competitive, depressing the performance of its rivals.

We address this challenge by adopting the approach of Bloom et al. (2013) who assess firm closeness in both dimensions simultaneously: in product space (using sales activity broken down by industry) and technology space (using patenting activity across technology fields). Bloom et al. (2013) distinguishes the effect of technology spillover from product market rivalry by basing its model specifications on the following skeleton:

$$Y = \beta_{tec}(D_{tec}G) + \beta_{sic}(D_{sic}G) + \epsilon$$

Simplifying to a single time period and to N firms, Y is an N x 1 matrix of a firm performance variable such as market value, knowledge, or productivity. The scalars β_{tec} and β_{sic}

are respectively the effect of technology spillovers and product market rivalry which are constant for all firms and are the quantities of interest in each regression. D_{tec} and D_{sic} are N x N distance matrices where the ijth cell represents the closeness between firm i and firm j on a zero-one scale (one meaning the firms are perfect neighbors). These matrices are determined prior to estimating the model. G is a column vector of length N where the ith element represents firm i's R&D expenditure.

By construction we set the diagonals of D_{tec} and D_{sic} equal to zero so that a firm's own R&D does not affect our measures of the potential for technology spillovers and the potential for business stealing. The N x 1 vector $(D_{tec}G)$ is identical to the $SPILLTEC$ variable of Bloom et al (2013). The ith row of $SPILLTEC$ is the sum of R&D of all firms other than firm i, with each weighted by its distance to firm i in *technology* space. Similarly $(D_{sic}G)$ is the N x 1 vector $SPILLSIC$ which may be described as the sum of R&D of all firms other than firm i each weighted by their distance to firm *i* in *product* space.

Factorizing yields

$$Y = (\beta_{tec}D_{tec} + \beta_{sic}D_{sic})G + \epsilon$$

In other words, the effect of a unit increase in the R&D of firm j (element j of the vector G) on the performance of firm i (element i of the vector Y) is a function of four things: the distance between the two firms in technology space (element ij of D_{tec}), a technology diffusion multiplier common to all firms (β_{tec}), the distance between the two firms in product market space (element ij of D_{sic}), and a product rivalry multiplier common to all firms (β_{sic}).

The strategy of Bloom et al. (2013) and this paper may be summarized in two steps. In step one we estimate the distance matrices D_{tec} and D_{sic}. In step two we multiply D_{tec} by G to yield $SPILLTEC$ and multiply D_{sic} by G to yield $SPILLSIC$. We regress firm performance metrics on SPILLTEC and SPILLSIC to yield β_{tec} and β_{sic}.

Estimating distance matrices in technology and product space

Following Bloom et al. (2013), we gauge how close two firms are in technology space by studying how frequently they patent in the same technology fields. For each firm, Bloom et al. (2013) estimates the average share of patents in each of 426 different technology fields from 1970-1999. This yields a time-invariant vector (T_i) of length 426 which describes patenting behavior from 1970-1999. We estimate the ij-th element of D_{tec} as the uncentered correlation between T_i and T_j (referred to in Bloom et al. (2013) as the Jaffe distance measure).

$$[D_{tec}]_{ij} = \frac{T_i T_j'}{(T_i T_i')^{1/2}\left(T_j T_j'\right)^{1/2}}$$

In an analogous fashion we gauge the proximity of firms in product space by studying how much revenue they earn in the same industries. Each firm's sales are separated into 597 different industry classes. Bloom et al. (2013) averages the share of each firm's sales in each of these classes over the period 1993-2001, yielding a time-invariant vector (S_i) of length 597 which describes sales from 1993-2001. We estimate the ij-th element of D_{sic} as the uncentered correlation between S_i and S_j.

$$[D_{sic}]_{ij} = \frac{S_i S_j'}{(S_i S_i')^{1/2}\left(S_j S_j'\right)^{1/2}}$$

In an extension to their models, Bloom et al. (2013) estimate D_{tec} and D_{sic} by developing their own Mahalanobis distance measure. The motivation of the Mahalanobis measure is to model spillovers as occurring between firms patenting (/earning revenue) in similar and not exclusively identical technology fields (/industry classes). We find that our results are in general robust to the choice of distance measure and so focus our discussion on the Jaffe distance

measure. For further discussion on the Mahalanobis distance measure, we refer readers to the Appendix.

Time-varying vs. time-invariant: A bias-variance tradeoff

Ideally patenting behavior and sales data would be perfectly representative of firm distances so that D_{tec} and D_{sic} could be precisely estimated each year. However, due to data limitations this is not always viable. On the technology side, between 1970 and 1997, data is available for an average of 82,000 patents per year. However, in 1998 we have data on only 33,780 patents and even fewer (1,560) in 1999. Sparsity is even more problematic on the product side. Performing our own fetch of the entire Compustat database, we find that a breakdown of sales by 4-digit industry code is available for a yearly average of 8,168 firms between 1999 and 2001. Compustat offers this breakdown for only 1,537 firms in 1998 and for less than 400 firms in 1995, 1996, and 1997 combined.

In order to yield more precise estimates of the placement of each firm in technology and product space, Bloom et al. (2013) averages patent and sales behavior over the entire period for which they have data. On the technology side, they assume that the distance between firms is time-invariant. On the product side, they assume that the distances are time-invariant and additionally, that these distances may be precisely estimated over a much smaller three-year window, 1999-2001.

By assuming distances are time-invariant, Bloom et al. (2013) achieve greater efficiency at the expense of a potential source of bias. Suppose a firm stops patenting in a technology field. Because distances are static, the firm will still be considered close to firms which patent in that field even though the mechanism of technology spillover has been extinguished. This model will

yield biased results in years which are not accurately represented by the firm's average placement in technology space.

The first and main contribution of this paper will be to relax the assumption that distances are static in technology space by creating a time-varying Jaffe measure. Data limitations prohibit us from doing the same in product space. However, the second contribution we make is to experiment with the construction of D_{tec} to roughly estimate the penalty to precision of only having used three years of data to construct D_{sic}.

Constructing a time-varying measure of firm distances in technology space

We extend the analysis of Bloom et al. (2013) by allowing D_{tec} to vary over time. Let T_{it} be a vector of length 426 with each element representing the share of patents made by firm i in year t that fall in each technology class. For any two firms i and j, we estimate the distance between the two firms in technology space using the moving average of patents in various technology fields from the previous, current, and subsequent years. We use the moving average in order to smooth the data, so that an unusual patenting year is not overly influential. Let $\bar{T}_{it}$ be a vector of length 426 with each element indicating the average share of patents by firm i in that technology field over the previous, current, and subsequent year.

$$\bar{T}_{it} = \frac{1}{3}\left(T_{i,t-1} + T_{i,t} + T_{i,t+1}\right)$$

Using this time-varying measure of patenting behavior, we construct a time-varying measure of the distance between firms i and j.

$$[D_{tec}]_{ijt} = \frac{\bar{T}_{it}\bar{T}'_{jt}}{(\bar{T}_{it}\bar{T}'_{it})^{1/2}\left(\bar{T}_{jt}\bar{T}'_{jt}\right)^{1/2}}$$

We demonstrate that, for any pair of firms i and j, $[D_{tec}]_{ijt}$ varies substantially over time. Note that because our estimates are based on fewer observations, they will exhibit more variability but less systematic bias.

Results

Do technology distances vary over time?

To test whether $[D_{tec}]_{ijt}$ is constant across years, we use a one-way ANOVA with random effects. For each possible combination of two firms, there is a series of values of $[D_{tec}]_{ijt}$ with one observation per year. For instance, we observe the distance between firm A and firm B in 1972, 1973, and so on through 1997. In this data structure, $[D_{tec}]_{ijt}$ is nested within the combination ij. Our goal is to compare the variance in $[D_{tec}]_{ijt}$ within a particular ij combination with the variance between different firm combinations. In the equation below μ_{ij} is a unique intercept for the distance between firm i and firm j, and ϵ_{ijt} is an idiosyncratic error term particular to the distance between firm i and firm j in year t.

$$[D_{tec}]_{ijt} = \mu_{ij} + \epsilon_{ijt}$$

This random-effects model assumes that the mean distance between any two firms, μ_{ij}, is randomly drawn from a normal distribution with mean μ and some variance σ_{μ}^2. We assume that the error term is normally distributed with mean zero and variance σ^2. We also assume that ϵ_{ijt} is independent of μ_{ij}, so the total variance of $[D_{tec}]_{ijt}$ can be written as $\sigma_u^2 + \sigma^2$.

$$\mu_{ij} \sim N(\mu, \sigma_{\mu}^2)$$

$$\epsilon_{ijt} \sim N(0, \sigma^2)$$

$$Var([D_{tec}]_{ijt}) = Var(\mu_{ij} + \epsilon_{ijt}) = Var(\mu_{ij}) + Var(\epsilon_{ijt}) = \sigma_{\mu}^2 + \sigma^2$$

The variance over time in the distance within a particular combination of firms is estimated by σ^2 and the variation between different combinations of firms is estimated by σ_μ^2. We summarize this information with the intraclass correlation (ICC), defined to be the proportion of variance that is between groups.

$$ICC = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma^2}$$

Bloom et al. (2013) assume that the distance between two firms does not vary over time. That translates into the assumption that $\sigma^2 = 0$ and the ICC is 1. If the estimated ICC is considerably less than 1, this will indicate that a substantial share of the total variance in technology distances is over time within firm combinations rather than between firm combinations.

Using the Jaffe distance measure, we find an intraclass correlation of 0.61. This indicates that 61% of the variance in technology distances is between different combinations of firms. Bloom et al. (2013) assume that the distance between any two firms does not vary over time, so they ignore 39% of the total variance in technology distances which is variation within firms over time. Although the assumption of time invariance is a reasonable one for the goal of precision, the large amount of variance within firms over time suggests that estimates ignoring this variance may be biased[1]. Because the ICC is far from 1, the authors may have ignored an important source of variation when they assumed that technology distances do not vary over time.

[1] Results using the Mahalanobis distance measure are similar. The Mahalanobis intraclass correlation is 0.73, indicating that 27% of the total variance in the Mahalanobis technology distance is within particular firm combinations over time.

Are the conclusions robust to time-varying technology measures?

In this section, we estimate the effects of technology spillovers using the time-varying technology distance measure. We replicated all of the models from tables 3-6 of Bloom et al. (2013) with and without allowing technology distances to vary over time. Results from all of these models were similar and are available in the appendix. We focus the discussion here on the models for Tobin's Q, which correspond to Table 3 in the original paper.

The dependent variable in these models is the log of Tobin's Q, defined to be market value divided by the firm's stock of fixed capital. This dependent variable captures the value added to a firm by technology spillovers or taken away due to market rivalry. We first replicate the OLS and fixed-effects models presented in Bloom et al. (2013), and then we estimate an additional model following their fixed-effects specification but allowing distances in technology space to vary over time.

Results are presented in Table 1. The sample size is smaller by about 38% in the model using the time-varying distance measure since this model excludes firms which did not patent in the year prior to the observation, the year of the observation, and the year after the observation. This exclusion is empirically necessary since we could not calculate placement in technology space for these firms, and it is reasonable since we would not expect firms which are not currently patenting to benefit from technology spillovers. The sample in Bloom et al. (2013) is already restricted to firms that patent sometime between 1970 and 1999, so our additional restriction is conceptually in agreement with the prior paper.

Model 1 is the OLS specification presented in Table 3, Column 1 of the original paper. All predictors are lagged by one year in order to ensure correct temporal ordering for a causal claim. Following Bloom et al. (2013), the model controls for a sixth-order polynomial for the

firm's own R&D stock divided by its capital stock, a dummy variable indicating firms with zero R&D stock, industry sales, and a series of year dummy variables. Net of these controls, a unit increase in the pool of R&D spending by technology neighbors ($SPILLTEC = D_{tec}G$) is associated with reduced market value while a unit increase in the pool of R&D spending by market rivals ($SPILLSIC = D_{sic}G$) is associated with increased market value. The authors expected technology spillovers to be beneficial and market rivalry to be detrimental, so these results are the opposite of what the authors predicted. However, the effects change sign in the model that includes firm fixed-effects (Model 2). The change suggests that time-invariant omitted variables confounded the original relationship. The authors take the results of Model 2 as evidence in favor of the theory that firms benefit when their placement in technology space is near other firms that invest heavily in R&D.

In Model 3, we relax the assumption that firms' placement in technology space does not vary over time. When this assumption is relaxed, the potential for technology spillovers is no longer an important predictor of market value ($\beta = -0.001$). The authors assumed placement in technology space to be constant over time in order to increase precision. The assumption seemed reasonable, but our model indicates that the results hinge on this decision.[2]

[2] We find similar results for models predicting the log of sales and the log of R&D investment divided by sales. These models, available in the appendix, show almost no effect of technology spillovers when distance in technology space is allowed to vary over time, regardless of whether the Jaffe distance measure or the Mahalanobis distance measure is used. We do find effects on citation-weighted patents similar to those presented by Bloom et al. (2013), but we think those results are driven by the choice of a method for estimating the negative binomial model with fixed effects. See the appendix for a discussion.

Omitted variable bias

Why might the fixed-effects results change so much when technology distances are allowed to change over time? Firm fixed-effects reduce bias by accounting for omitted variables which are constant within firms. However, they rely entirely on variation within firms over time. Change over time in $SPILLTEC$ in Model 2 is entirely driven by changes in the R&D spending of nearby firms, while change over time in $SPILLTEC$ in Model 3 also includes changes in which firms are nearby. Therefore, these two models estimate slightly different quantities of interest. We argue that the fixed-effects results from Bloom et al. (2013) using time-invariant technology distances may be particularly prone to omitted variable bias. To the extent that there are unobserved fluctuations in the economy for a particular technology space, these fluctuations may affect both the market value of the firm of interest and the R&D spending of nearby firms. A model that assumes placement in technology space does not change and thus relies solely on changes in R&D spending will be highly prone to bias due to these omitted variables[3].

Our models that allow for changes in technology space over time may also be limited by omitted variable bias. Unobserved variables affect firms' placement in technology space, their

[3] Bloom et al. (2013) address the possibility of omitted variables in alternative models by using tax incentives as an instrument for the R&D spending of nearby firms. Although their IV models support the conclusion that technology spillovers are beneficial, we find a correlation of 0.98 between the instrument for $SPILLTEC$ and the non-instrumented value of $SPILLTEC$, using the data and code provided on Bloom's website. IV analysis assumes that omitted variables affect $SPILLTEC$ but do not affect the instrument. We find this assumption questionable because $SPILLTEC$ is so highly correlated with the instrument. Therefore, we focus here on the models that do not instrument for $SPILLTEC$.

market value, and the R&D spending of other firms. Therefore, our results do not conclusively answer the question of whether firms benefit from technology spillovers. To the extent to which omitted variables predict both the potential for technology spillovers and the success of a firm, these omitted variables should bias our coefficients upward. It is therefore interesting that our results yield effect sizes near zero, implying that the NBER dataset may not lend as much evidence to the existence of technology spillovers as Bloom et al. (2013) suggests.

Attenuation due to measurement error

Our findings may be limited due to measurement error. It could be the case that firms' placements in technology space are actually constant, and the variability we found over time is simply due to measurement error between years. In this case, we would have measurement error in $SPILLTEC$, which would attenuate the coefficient on $SPILLTEC$ toward zero in the models that allow D_{tec} to vary over time. If measurement error were severe, it might explain our lack of a technology spillover effect.

Although we cannot entirely rule out measurement error, we can lend suggestive evidence as to whether it may be responsible for the results. In the model allowing time-varying technology distances, we found a coefficient on $SPILLTEC$ of -0.001. If technology spillovers actually have a positive effect, then it would take an enormous amount of measurement error in $SPILLTEC$ to yield such a small coefficient. Presumably, a variable with this level of measurement error would have a very high variance. However, we find that the variance of our $SPILLTEC$ measure is 1.26, compared with a variance of 1.10 for the measure used in Bloom et al. (2013). It is still possible that our measure includes more measurement error that is not captured by the overall variance, but the similar variances provide some reassurance that measurement error is not so extreme as to entirely drive our results.

Measurement error in product market distances

In the previous discussion, we have focused on the measurement of placement in technology space and the coefficient β_{tec} because better data is available for technology distances than for product market distances. However, conclusions about β_{tec} rely heavily on conclusions about β_{sic} since technology spillovers and product market rivalry are competing effects. The relative size of β_{tec} and β_{sic} provides policymakers with an indication of how close two firms must be in technology space for the positive effect of technology diffusion to overpower the negative effect of their rivalry in product space. Recall that a unit increase in the R&D expenditure of firm j has an effect of $\beta_{tec}[D_{tec}]_{ij} + \beta_{sic}[D_{sic}]_{ij}$ on the market value of firm i and suppose β_{tec} is twice the absolute value of β_{sic}. Then a unit increase in the R&D spending of all firms more than half as close to firm i in technology space as they are in product space will net a positive effect on the market value of firm i. In other words, the relative size of β_{tec} and β_{sic} furnishes policymakers with a set of firms whose R&D nets a positive effect on their peers and so are good candidates for targeted subsidies.

Whereas Bloom et al. (2013) estimate D_{tec} using almost thirty years of data, the limitations of the Compustat database constrain them to estimating D_{sic} using only three full years of data from 1999-2001. We expect more measurement error and therefore downward bias in β_{sic} than in β_{tec} simply for having calculated D_{sic} using a smaller set of observations. Bloom et al. (2013) find β_{tec} to be 4.6 times larger than β_{sic}. In this section we gauge how much of this spread is due to extra measurement error in β_{sic}.

To do so we assume that a year of patent data is as informative about technology space as a year of revenue data is about product space. Ideally, we would use the same years to calculate both technology distances and product market distances. We are unable to calculate D_{sic} using

data from 1970, however we are able to deliberately impair our estimates of D_{tec} to be comparable to the estimates of D_{sic}. We engage in the following thought experiment.

Just as D_{sic} is measured using three years of data we deliberately occlude all but three years of patent data to generate D_{tec} and then proceed to estimate β_{tec}. We roll this three year window forward to yield Figure 1, a plot of β_{tec} against the three years of underlying patent data used to estimate D_{tec}. Also plotted is a baseline of 0.381 derived from averaging patent data across 1970-1999 in Bloom et al. (2013).

In line with expectations we see that the occluded estimates of β_{tec} are consistently lower than the full-sample estimate generated in Bloom et al. (2013). On average β_{tec} is 86% smaller than the full-sample estimate. If β_{sic} suffers to as great a degree of measurement error as did our experimental β_{tec} estimates, then β_{sic} could be as high as -0.592 as opposed to –0.083 (calculated by scaling -0.083 up so that it constitutes 14% of -0.592). This adjustment to β_{sic} is sufficiently large as to render β_{sic} greater than β_{tec}, greatly narrowing the set of firms that are good candidates for targeted subsidies. In particular, after rescaling β_{sic} for measurement error we find that firms must be closer in technology space than they are in product space in order for their R&D spending to net a positive effect on their peers.

Conclusion

Bloom et al. (2013) argue that firms benefit from R&D spending by firms in their technology space but are harmed by R&D spending by firms that share their same product market. This paper has expanded on this work by discussing potential problems in the measurement of technology spillovers and product market rivalry. Our first contribution is to investigate alternative specifications for the potential for technology spillovers. Bloom et al. (2013) assume that the distance between any two firms in technology space is constant over time.

Relaxing this assumption, we find no evidence of a positive effect of technology spillovers on market value. Since the measure used in Bloom et al. (2013) is highly dependent on changes in R&D spending of technology neighbors, we suggest that their results may be driven by bias from omitted variables related to changes in the economic conditions in a particular technology niche.

Our second contribution is to discuss potential problems with the measurement of product market rivalry. While the potential for technology spillovers is measured using 30 years of data, product market rivalry is measured using only 3 years of data. The lack of precision in the estimate of product market rivalry may attenuate the effect of product market rivalry, which could yield artificially high estimated returns to technology spillovers since the two are opposing forces. When we use a similar method to estimate technology spillovers using only 3 years of data, we find a much larger role for product market rivalry and a smaller role for technology spillovers.

Taken together, these findings suggest that Bloom et al. (2013) may have underestimated business stealing effects from product market rivalry and overestimated the social returns to R&D through technology spillovers. Before incentivizing R&D spending, policymakers should conduct more research to identify the precise conditions under which R&D leads to beneficial technology spillovers.

References

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying technology spillovers and product market rivalry." *Econometrica* 81(4):1347-1393.

Table 1. Effects of technology and product market spillovers, using Jaffe distance measure

	Time-invariant tech distance (Bloom et al. 2013)		Time-varying tech distance
	(1)	(2)	(3)
Firm fixed-effects?	No	Yes	Yes
Log(Tobin's Q):			
ln(SPILLTEC$_{t-1}$)	-0.064	0.337	-0.001
	(0.013)	(0.108)	(0.017)
ln(SPILLSIC $_{t-1}$)	0.053	-0.085	0.041
	(0.007)	(0.032)	(0.049)
No. of observations	9,944	9,944	6,209

Note: Models 1 and 2 are direct replications of Bloom et al. (2013). Model 3 uses the same specification but with an alternate measure of $SPILLTEC$ that allows the distance between firms in technology space to change over time. Tobin's Q is defined as the ratio of the firm's market value to the firm's stock of fixed capital. Models control for a sixth-order polynomial for the firm's own R&D stock divided by its capital stock, a dummy variable indicating firms with zero R&D stock, industry sales, and a series of year dummy variables. The corresponding models are presented by Bloom et al. (2013) in Table 3. All OLS standard errors are corrected for autocorrelation of order 1 using a Newey-West correction. Negative binomial standard errors are clustered to account for non-independence within firms.

Figure 1. Measurement error in occluded β_{tec} estimates using a 3-year rolling window

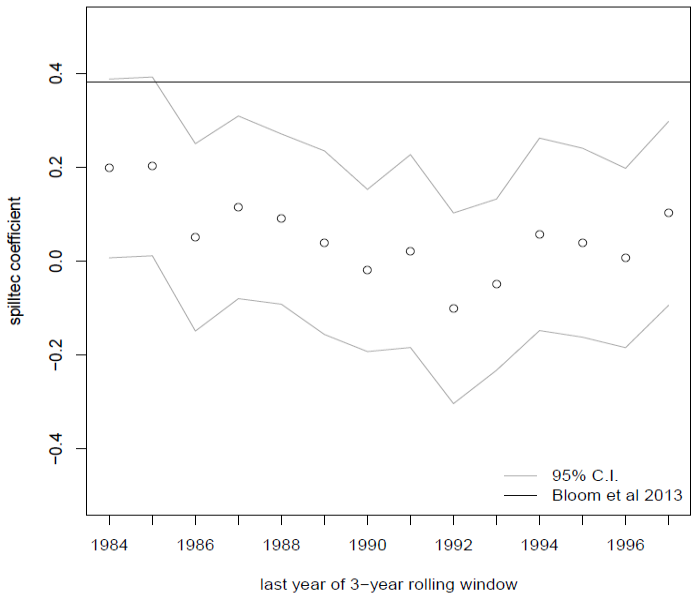

Note: Consider the first data point. Using a distance matrix (D_{tec}) constructed on patent data between 1982-1984, we estimate β_{tec} to be 0.20. The shortfall between this occluded estimate of β_{tec} and the full-sample estimate of 0.381 provided in Bloom et al. (2013) is used to proxy the measurement error in β_{sic} from only having used three years of data to construct D_{sic}.

Appendix: Alternative dependent variables

In addition to the models predicting Tobin's Q, we estimated additional models replicating the results from tables 4-6 of Bloom et al. (2013). The results generally agreed with those from the Tobin's Q model presented in the main text, with the exception of the negative binomial models predicting citation-weighted patent counts. We discuss these results below and present the coefficient estimates in Table A1.

Following Bloom et al. (2013), we weight by citations in order to account for the higher value of patents which receive more citations. Models 2 and 3 estimate fixed effects using a "pre-mean scaling approach" (Blundell, Griffith, and Van Reenen 1999), which we use in order to follow the method used in Bloom et al. (2013). Our results for these models agree with the previously published results; firms produce more patents if they are have a higher potential for technology spillovers, and this effect is robust to the use of time-invariant or time-varying measures of technology distance.

We suspect that the cite-weighted patent results may agree with those published by Bloom et al. (2013) because the method used to estimate the fixed effects is inadequate. The pre-mean scaling approach assumes that one can estimate each firm's individual fixed-effect, or time-invariant propensity to patent, based on a long series of observations prior to the sample (Blundell, Griffith, and Van Reenen 1999). If this method yields inaccurate estimates of the firm fixed-effects, then it does not resolve the problem of time-invariant omitted variables. Since the citation-weighted patents is the only dependent variable for which our model yields a large positive effect of technology spillovers, we suspect that these results are driven by omitted variable bias.

The next set of models regress the log of sales on the potential for technology spillovers, the potential for business stealing, and control variables. The OLS results using time-varying technology distances (Model 3) agree with the OLS results presented in Bloom et al. (2013) (Model 1), but the fixed-effects models are different. As in the models predicting Tobin's Q, the time-invariant technology distance measure used in Bloom et al. (2013) yields a positive effect of technology spillovers on sales, but the time-varying technology distance measure shows no effect of technology spillovers on sales.

The final set of models in Table 1 predict the log of R&D investment divided by sales. These models test whether firms invest more in R&D when they have greater potential to benefit from technology spillovers, possibly because the returns to R&D could be higher if a firm is beneficially located in technology space. In the fixed-effects specification, Bloom et al. (2013) find that firms spend more on R&D if nearby firms are also investing in R&D (Model 2), though the effect is not statistically significant ($\beta = 0.100, SE = 0.076$). When we allow the distance between firms to vary with time (Model 4), we also estimate a non-significant positive effect of technology spillovers, but our estimated effect size is much smaller ($\beta = 0.013, SE = 0.014$).

In Table A2, we present similar results using the Mahalanobis distance measure. The results are essentially the same as those using the Jaffe distance measure. The fixed-effects models allowing technology distance to vary over time do not demonstrate an effect of technology spillovers, with the exception of the model predicting citation-weighted patents.

Table A1. Effects of technology and product market spillovers, using Jaffe distance measure

	Time-invariant tech distance (Bloom et al. 2013)		Time-varying tech distance
	(1)	(2)	(3)
Firm fixed-effects?	No	Yes	Yes
Log(Tobin's Q):			
$\ln(SPILLTEC_{t-1})$	-0.064	0.337	-0.001
	(0.013)	(0.108)	(0.017)
$\ln(SPILLSIC_{t-1})$	0.053	-0.085	0.041
	(0.007)	(0.032)	(0.049)
No. of observations	9,944	9,944	6,209
Cite-weighted patents:			
$\ln(SPILLTEC_{t-1})$	0.518	0.468	0.308
	(0.078)	(0.068)	(0.040)
$\ln(SPILLSIC_{t-1})$	0.045	0.056	0.009
	(0.035)	(0.032)	(0.025)
No. of observations	9,023	9,023	5,544
Log(Sales):			
$\ln(SPILLTEC_{t-1})$	-0.022	0.191	0.001
	(0.009)	(0.046)	(0.008)
$\ln(SPILLSIC_{t-1})$	-0.016	-0.005	0.010
	(0.004)	(0.011)	(0.023)
No. of observations	9,935	9,935	6,226
Log(R&D / Sales):			
$\ln(SPILLTEC_{t-1})$	0.079	0.100	0.013
	(0.018)	(0.076)	(0.014)
$\ln(SPILLSIC_{t-1})$	0.374	0.083	0.106
	(0.013)	(0.034)	(0.034)
No. of observations	8,579	8,579	6,509

Note: The first two columns are a direct replication of Bloom et al. (2013). The second set of columns use the same specification, but they use an alternate measure of $SPILLTEC$ that allows the distance between firms in technology space to change over time. Tobin's Q is defined as the ratio of the firm's market value to the firm's stock of fixed capital. The corresponding models in Bloom et al. (2013) are Table 3 for Tobin's Q, Table 4 for cite-weighted patents, Table 5 for sales, and Table 6 for R&D. For details on model specification, see the notes for those tables in the original paper. All models are OLS models except for the model of citation-weighted patents, which is a negative binomial model. All OLS standard errors are corrected for autocorrelation of order 1 using a Newey-West correction. Negative binomial standard errors are clustered to account for non-independence within firms.

Table A2. Effects of technology and product market spillovers, using Mahalanobis distance measure

	Time-invariant tech distance (Bloom et al. 2013)	Time-varying tech distance
	(1)	(2)
Firm fixed-effects?	Yes	Yes
Log(Tobin's Q):		
$\ln(SPILLTECH_{t-1})$	0.759	0.012
	(0.137)	(0.018)
$\ln(SPILLSIC_{t-1})$	-0.139	0.118
	(0.050)	(0.072)
No. of observations	9,944	6,209
Cite-weighted patents:		
$\ln(SPILLTECH_{t-1})$	0.583	0.392
	(0.080)	(0.049)
$\ln(SPILLSIC_{t-1})$	0.079	0.016
	(0.047)	(0.039)
No. of observations	9,023	5,544
Log(Sales):		
$\ln(SPILLTECH_{t-1})$	0.264	0.005
	(0.064)	(0.008)
$\ln(SPILLSIC_{t-1})$	-0.007	0.023
	(0.021)	(0.032)
No. of observations	9,935	6,226
Log(R&D / Sales):		
$\ln(SPILLTECH_{t-1})$	-0.176	0.012
	(0.101)	(0.015)
$\ln(SPILLSIC_{t-1})$	0.224	0.192
	(0.048)	(0.047)
No. of observations	8,579	6,509

Note: The first two columns are a direct replication of Bloom et al. (2013). The point estimates are not exactly the same as in the original paper, but they are close enough to be substantively the same. Although the authors do not present the models without fixed-effects shown here in the first column, we use the same specification and variables from their fixed-effects models but omit the fixed effects in this column. The second set of columns use the same specification, but they use an alternate measure of $SPILLTEC$ that allows the distance between firms in technology space to change over time. Tobin's Q is defined as the ratio of the firm's market value to the firm's stock of fixed capital. The corresponding models in Bloom et al. (2013) are Table 3 for

Tobin's Q, Table 4 for cite-weighted patents, Table 5 for sales, and Table 6 for R&D. For details on model specification, see the notes for those tables in the original paper. All models are OLS models except for the model of citation-weighted patents, which is a negative binomial model. All OLS standard errors are corrected for autocorrelation of order 1 using a Newey-West correction. Negative binomial standard errors are clustered to account for non-independence within firms.