Identifying Technology Spillovers and Product Market Rivalry

by Nicholas Bloom, Mark Schankerman, and John Van Reenen

in *Econometrica* 81(4): 1347-1393. (2013)

Research Digest by Matthew Clancy

Empirical Contribution

This is a long paper that discusses many topics. This research digest will zero in on the unambiguous theoretical predictions and the methods used to test them.

Major Predictions

This paper is about the impact of research and design (R&D) by firms on each other. R&D can impact other firms through two major channels, which the authors call the *technology spillover* channel and the *product market rivalry* channel. To illustrate these channels, consider the case of three hypothetical firms: BikeCo, eScooter, and DroneCo.

BikeCo manufactures bikes. eScooter manufactures rechargeable scooters. DroneCo manufactures small flying drones. BikeCo and eScooter both compete for the same customer base of people interested in convenient urban transportation, but use completely different technological approaches (bicycles versus electric scooter). Meanwhile, eScooter and DroneCo sell to completely different customer bases but rely on similar underling technologies (better batteries and more efficient engines). Suppose each of these companies spends $1 million per year on R&D. If either BikeCo or DroneCo doubles its spending on R&D, it will have an impact on eScooter, but the impact will differ significantly depending on which firm doubles its spending.

If DroneCo doubles its R&D, the additional advances it makes in battery technology and more efficient motors can be readily adapted by eScooter to improve its own product line. If it continues to spend $1 million per year on R&D, this spending will generate more innovation than before, leading to more profit. Or, eScooter may decide it can reduce its R&D spending and free-ride on the efforts of DroneCo, obtaining a similar level of innovation with fewer R&D costs, again raising its profit. The impact of DroneCo's R&D on eScooter represents a *technology* spillover, where R&D by a rival firm enhances R&D capacity and raises profit.

However, if BikeCo doubles its R&D, the ramifications for eScooter are much less positive. Suppose BikeCo does R&D on material science to develop cheaper, sturdier and lighter bike frames. BikeCo's advances in material science have little application for scooters, and so eScooter's R&D capacity is unchanged. But it now faces stiffer competition from better bikes, and so it may need to reduce its profit margin or perhaps spend more on R&D itself. Either way, its profitability is negatively impacted. In this case, the impact of BikeCo's R&D on eScooter represents a *product market rivalry* spillover, where R&D by a rival firm has no impact on R&D capacity but reduces profit.

The preceding is an extreme example, but the more general point is that R&D by rivals has positive and negative effects. The positive effect is via the technology spillover, where firms learn from each other and borrow/adapt their ideas. This effect is strongest when firms operate in similar technological fields. The negative effect is via the product market rivalry spillover, where rival R&D by rival firms makes for tougher competition. This effect is strongest when firms compete in similar markets. In most cases firms face a mix of technology and product market spillovers from their rivals.

This paper works hard to separately identify and measure these two channels. Without separately identifying these channels, the total impact of rival firms R&D will be understated. Specifically, the authors develop a measure of technology spillovers from rival firms that they call *SPILLTECH*, and a measure of product market rivalry spillovers from rival firms that they call *SPILLSIC* (so-called because SIC codes are a common way of defining different industries). They compute these measures for a large sample of firms over two decades, and then assess whether they are correlated with innovation and profitability in the manner predicted. Specifically, they predict, holding all else constant:

1. *SPILLTECH* is positively correlated with a patent-based measure of innovation and a productivity-based measure of innovation
2. *SPILLSIC* is uncorrelated with a patent-based measure of innovation and a productivity-based measure of innovation
3. *SPILLTECH* is positively correlated with profitability
4. *SPILLSIC* is negatively correlated with profitability

Measuring Technological Spillovers

A key contribution of the paper is to derive a way to measure the extent of technology spillovers from rival firms. To motivate their measure of spillovers, let’s imagine a highly simplified and unrealistic model of knowledge spillovers. Suppose eScooter employes 10 scientists and DroneCo employes 5. These scientists all attend an annual meeting where they will meet all of their counterparts in other firms. The 10 scientists of eScooter will meet each of the 5 scientists in DroneCo, for a total of 50 meetings. At each of these meetings, there is some probability that a DroneCo scientist will divulge a piece of information useful to the eScooter scientist. We call this a “knowledge spillover.” To begin, assume these information transfers are only useful if the two scientists are working the same technology field (we’ll relax this assumption in a minute). How much useful information will DroneCo transfer to eScooter?

It depends on the share of scientists working in the same field. Suppose the 10 scientists of eScooter work in 5 different fields (2 scientists per field), and the 5 scientists of DroneCo work in the same 5 fields (1 scientist per field). Then eScooter scientists working in field *A* (say, battery charging technology) only get useful information in the meetings they have with the single scientist in DroneCo working in that field. The rest of the meetings are a waste of time. Since each of the 10 scientists in eScooter has a single useful meeting with scientists in DroneCo, the total spillovers from DroneCo to eScooter are 10.

Note this is a symmetric exchange. The 5 DroneCo scientists each meet with two eScooter scientists in their field and learn two useful pieces of information from them as well, so the total spillovers from eScooter to DroneCo are also 10.

Now, this isn't realistic, but it captures some useful intuitions. If you have a rival firm that does a lot of R&D, you stand to learn a lot from them, especially if they tend to work in the same kinds of areas as your own firm. Alternatively, if you do a lot of R&D yourself, you have more capacity to absorb knowledge from rivals.

Bloom, Schankerman and Van Reenen don't actually have data on the number of scientists employed by each firm, broken down by their fields though. Instead, they have data on each firm's total R&D spending and patents. Fortunately, these patents are assigned to 426 different technology classes by the US Patent and Trademark Office. For example, class 320 corresponds to "Electricity: battery or capacitor charging or discharging." Bloom, Schankerman, and Van Reenen assume the number of researchers working in the field is proportional to R&D spending (technically, they construct measures of R&D that include partial contributions from previous years) multiplied by the *share* of patents in the field. So if eScooter spends $1 million per year on R&D and 40% of it's patents belong to class 320 (batteries), then the assume the company spends $400,000 per year on battery research.

Note that the number of information transfers in the illustrative example is given by multiplying the number of scientists in every field by the number working in the same field at the other firm, and then adding them up. Since Bloom, Schankerman, and Van Reenen are now using estimates of R&D spending in each field, if they were to calculate spillovers from DroneCo to eScooter, they would take an estimate of R&D eScooter does in every field and multiply it by their estimate of R&D DroneCo does in the same field. They do this for all 426 fields and then add up the results to obtain the total spillover.

But there is a problem with this measure. The assumption that knowledge transfer can only occur between the same fields is unrealistic; technology fields frequently borrow ideas from each other. For example, maybe the next advance in battery technology will come advances in chemical processing. So the authors want there to be a non-zero probability that knowledge transfer occurs between scientists working in different fields, even if this probability is probably lower than it would be if they are in the same field.

To derive this probability, the authors assume firms do not choose different research fields at random. Firms can benefit from internal knowledge transfers across fields too. Since these are a good thing for the firm, firms have an incentive to do research on the fields most likely to generate knowledge transfers between the different scientists working for them. The authors assume that fields that frequently appear together in the same firms are more likely to have knowledge flows between them. Specifically, to find the probability a meeting between scientists in field *A* and *B* have a knowledge transfer, they find the share of patents in every firm that belong to fields *A* and *B* and take the (uncentered) correlation between these two shares across firms. So, for example, if knowing the share of patents in class 320 (batteries) for a firm perfectly predicts the share of their patents in class 324 (measuring electricity), then the correlation is 1 and a knowledge transfer between the fields is guaranteed. Conversely, if knowing the share of patents in class 320 (batteries) for a firm confers no predictive power for class 300 (brush, broom, and mop-making), then there is zero correlation between the classes and the probability of a knowledge transfer is zero.

So the actual measure of spillovers between any two firms *A* and *B* is derived by adding up the product of firm *A*'s R&D in each field, firm *B*'s R&D in each field, and the correlation between fields. For example, to find the technology spillover from DroneCo to eScooter, the authors start with the first technology class (002 - apparel). They multiply their estimate of eScooters's R&D in class 002 by their estimate of DroneCo's R&D in class 002. Next, they multiply eScooter's R&D in class 002 by DroneCo's R&D in the next class (004 – baths, closets, sinks, spittoons), multiplied by the correlation between classes 002 and 004. They repeat this for all 426 fields to find the total spillovers from DroneCo to eScooter's class 002 R&D. Then they do the same for the eScooter class 004 R&D, then class 005, and so on until they've done all 426 fields for eScooter. They add up all these spillovers to compute the total extent of spillovers from DroneCo to eScooter.

Finally, note that DroneCo is not the only company doing research relevant to eScooter. What matters is the total amount of spillovers from *all* firms. So the main measure of interest for Bloom, Schankerman and Van Reenen is actually the sum of technology spillovers from all other firms. They call this variable *SPILLTECH*, and they compute it for every firm in every year in their dataset. When it's large, that means there are a lot of firms doing a lot of R&D in fields that are closely related to your firm. In plain English, firms have lot of opportunity to learn from their rivals.

Measuring Product Market Rivalry Spillovers

Bloom, Schankerman and Van Reenen follow a similar approach to compute the extent of market rivalry spillovers. Once again they have measures of R&D by each firm but now they want to estimate how much of that R&D is used to improve products in different markets. To estimate the share of R&D that is applicable to different product markets, the authors use the share of sales going to 597 different industries, as defined by the four digit SIC classification system. For example, the majority of eScooter and BikeCo's sales are in industry 3571 (which includes both bicycles and motorized scooters).

As in the previous case, it's possible that R&D might be important for multiple sectors. For example, maybe 100% of BikeCo's patents are in class 420 (alloys or metallic compositions), but it applies this knowledge to its main business of bikes, but also to a side-business of manufacturing metal bike tools. To allow for this, they perform the same calculations as before to find the correlation between different industry sales. What that means is that, for example, to compute the correlation between SIC industry 3571 (bikes, scooters, and more) and industry 3423 (handtools), they find every firm's share of sales in industry 3571 and share of sales in industry 3423 and compute the correlation between the two. When the share of sales in industry 3571 is a perfect predictor of the share of sales in 3423, then the two are perfectly correlated and obtain a correlation coefficient of 1.

To compute the product market rivalry spillover between some firm *A* and firm *B*, the authors start with the first SIC four-digit industry in their sample, say industry 2013 (sausage and other meat products). They take the share of sales from firm *A* in industry 2013, the share of sales from firm *B* in 2013 and multiply them together with their estimate of firm *B*'s R&D (there is an additional step involving division by a constant that I am omitting for simplicity). This is their measure of how firm *B*'s R&D to improve their industry 2013 products impacts firm *A* via its products in industry 2013. Then they move on to the next industry in firm *B*, say industry 2015 (poultry processing). They multiply the share of sales from firm *A* in industry 2013, the share of sales from firm *B* in industry 2015, the *correlation* of sales shares between 2013 and 2015, and firm *B*'s R&D. This is their measure of how firm *B*'s R&D in industry 2015 impacts firm *A* via its industry 2013 products. They repeat this for all 597 industries to compute the total impact of firm *B*'s R&D on firm *A*'s products in industry 2013. Then they do it all again for firm *A*'s sales in industry 2015, then industry 2021, and so on for all 597 industries. Adding it all together, they obtain a measure of R&D by firm *B* that impacts firm *A* via product market rivalry.

However, as in the other case, firm *A* doesn't just care about firm *B*. Instead, it's strategy will be affected by the R&D of all other firms operating in markets competing with firm *A*. So Bloom, Schankerman and Van Reenen add up the product market rivalry spillover from each firm to construct a variable they call *SPILLSIC*. When a firm's *SPILLSIC* is large, it means other firms operating in similar product markets are doing a lot of R&D; it implies rivals are stepping up their game.

Empirical Results

Now that they have a way to measure spillovers via the technology and product market channels, Bloom, Schankerman, and Van Reenen are in a position to test their predictions. They have a sample of annual observations over 1980-2001 for 715 publicly traded, patenting, US firms.

To test prediction 1 and 2, they measure innovation in two ways. First, they use a firm's patents as a proxy for innovation, weighting each patent by the number of citations it receives in the years after being granted. This form of weighting is a standard way to account for the vast differences in the importance of different patented inventions. The authors run a regression with *SPILLTECH*, *SPILLSIC*, and the firm's own R&D as explanatory variables for the number of citation-weighted patent applications by the firm per year. As they predict, there is a positive correlation between the number of patents and *SPILLTECH*: the more R&D by rival firms in similar technology classes, the more patents, after holding constant a firm's own R&D and *SPILLSIC*. There is also a positive correlation between *SPILLSIC* and patents, but it is so small that they cannot reject the hypothesis that it is equal to zero. That is, holding constant a firm's R&D and spillovers from similar technology classes, there is no independent impact from R&D by firms operating in the same product market.

Their second way of measuring innovation is based on the productivity of firms. For a given level of labor and capital (which they have annual data on), how much quality-adjusted output can the firm produce? To measure this, the authors need a measure of quality-adjusted output: they use sales divided by an industry-specific price index to control for inflation over time. They then include measures of capital, labor and own-firm R&D as explanatory variables. The intuition is that firms that make more revenues, after adjusting for inflation, with the same quantity of inputs must be producing better outputs or are using inputs more efficiently to produce more outputs. It is reasonable to believe R&D (both by the firm itself and rivals operating in similar technology categories) could improve product quality and production processes. While this is not a perfect measure, it does not have the same biases as using patents to measure innovation, and so provides a useful second test. Regressing this measure of sales on *SPILLTECH*, *SPILLSIC*, own firm R&D, labor, and capital, they obtain results as predicted. Firms enjoying more R&D by rivals in the same technology field (higher *SPILLTECH*) have higher output, holding constant everything else. Firms experiencing more R&D by rivals in the same product market (higher *SPILLSIC*) see no statistically significant impact, holding everything else constant.

To test prediction 3 and 4, the authors need a measure of firm value. They use the market value of the firm, since this embodies the markets assessment of the value of all predicted future profits of the company. Adjusting market value by the value of non-R&D assets and adding in firm-specific fixed effects they find value is positively correlated with *SPILLTECH* and negatively correlated with *SPILLSIC*, as they predict. A 10% increase in *SPILLTECH* is associated with a 3.8%-10.8% increase in market value (depending on the exact model), and a 10% increase in *SPILLSIC* is associated with a 0.8%-2.4% reduction in market value (depending on the model). Both results are statistically distinguishable from zero. Both effects are as predicted: R&D by firms in technologically similar areas raises value, while R&D by firms operating in similar product markets lowers value.

Additional Contributions

The paper makes a number of additional contributions that I only briefly mention here.

First, one concern with the above exercise may be that omitted variables are driving the result. The problem is that the decision about how much R&D to do responds to the outlook for R&D. When the outlook is good, many firms will simultaneously raise their R&D, and we might attribute the subsequent increase in innovation to spillovers instead of the improved R&D outlook.

Ideally, Bloom, Schankerman, and Van Reenen would address this by taking the choice about R&D out of the hands of the firms. They would randomly assign some firms to do more R&D and others to do less, and then run their regressions. Naturally, such an experiment is completely infeasible as a practical matter.

Instead, they exploit variation in tax incentives to obtain a similar kind of result. The “cost” of R&D varies across firms due to differential tax policy. For example, firms operating in different states face different R&D tax incentives. Since firms are rational, we assume they will do more R&D when taxes make it cheaper, and less when tax policy makes it more expensive. While the decision about where to conduct R&D is not random, the authors argue *changes* in tax policy are all but random for any specific firm. Changes in tax policy will, in turn lead to changes in R&D. If the tax policy changes are plausibly random, then the changes in R&D they induce can also be thought of as plausibly random. Some firms end up increasing their R&D because they were lucky enough to operate in states that decided to increase R&D tax credits, others end up decreasing R&D because they were unlucky and operate in states that ended R&D tax credit programs. The authors use an approach called instrumental variables to isolate these tax-driven R&D changes, which are plausibly random, and run their regressions with these changes. This exercise does not substantively changing their conclusions.

Second, the paper examines the robustness of their results to a host of alternative assumptions. They add in a geographic dimension, so that firms can enjoy more spillovers from rivals in the same broadly defined region. They explore alternative ways of measuring *SPILLTECH* and *SPILLSIC*. They use alternative data sources. They restrict their analysis to high-tech firms only. Again, they generally find their conclusions stand.

Third, the paper carefully develops a theoretical model to more fully flesh out the impact of spillovers than described here. In particular, the model discusses how spillovers impact the amount of R&D a firm should conduct, showing the results are ambiguous. The paper also discusses the strengths and weaknesses of different ways of measuring spillovers, arguing that while their preferred model has some attractive properties, there does not exist any measure which is obviously superior along every desirable criterion.

Lastly, the paper uses its regression results to estimate the social and private rate of return to R&D. As this paper makes clear, R&D decisions by one firm impact others. If firms do not take into account the impact of their R&D decisions on rivals, then the amount of R&D they choose to do will deviate from the socially optimal level. Specifically, R&D entails benefits to rival firms operating in the same technology field and costs to rival firms operating in the same product market (via a business stealing effect). Because Bloom, Schankerman, and Van Reenen have estimates of how much R&D affects the output of other firms, they can use them to estimate the actual return on R&D. They find the average social rate of return is 55% (meaning society as a whole enjoys a return of $1.55 for every dollar of R&D spent), and the average private rate of return is 21% (meaning individual firms earn $1.21 for every dollar of R&D spent).

Discussion

The fact that the social return on R&D is estimated to be more than twice the return enjoyed by the firm doing R&D underscores just how important spillovers are. Taking the estimates at face value, more than half the value of R&D comes from spillovers!

Moreover, the empirical exercise also underscores the importance of separating the technology spillover from the product market rivalry spillover, since the two effects operate differently and sometimes at cross purposes. In most cases, rival firms overlap in technology and product markets – think IBM and Apple – so that the total spillover is a mix of the two effects. If we attempted to measure the extent of spillovers with *only* a measure like *SPILLTECH* or *SPILLSIC* (but not both), our estimate would be too low. This is because the log of the two measures have a correlation of 0.4, and so in isolation their apparent impact on something like market value is a mix of the positive technology spillover and the negative product market rivalry spillover. In fact, when the authors run their regressions with either variable in isolation, they find smaller effects than when they allow them both to operate.

Data

The paper primarily relies on two data sources. It obtains data on firms from the U.S. Compustat dataset. This gives them firm-level accounting data such as R&D spending, sales (total and broken down by four digit industry code), market value, employment, and capital for all publicly traded US firms over 1980-2001. Because there is some debate about the accuracy of the Compustat data on sales broken down by industry, they check their results hold when they instead use an alternative datas source for this sales by industry, called BVD. To measure real output, they divide sales by industry price deflators taken from another paper (Bartelsman, Becker and Gray 2000) and the BEA four digit NAICS Shipment Price Deflators (only for years after 1996).

As noted above in passing, the authors do not actually use annual R&D to construct their spillover measures, but rather a measure of R&D that gives some value to prior R&D expenditure. This is a common practice in the economics of innovation. The intuition is that what really determines a firm's spillovers and innovative output is not so much the R&D they spend in any given year but the total knowledge the firm has accrued. The idea is that R&D generates knowledge, but then this knowledge sticks around and continues to be useful for a time in the future.

Specifically, they create a variable called an R&D stock. In the first year of a firm's existence, it's R&D stock is just equal to it's R&D in the first year. But for every year after, the R&D stock is equal to 85% of its level in the previous period, *plus* any R&D conducted in the current period. For example, if the R&D stock is $100 million in 2000, and a firm spends $20 million on R&D in 2001, then in 2001 it will be equal to 0.85 x $100mn + $20mn = $105 million. The argument for why the R&D stock decays by 15% each period is that knowledge generated by R&D becomes less useful for generating new innovations over time, as the current state of the art advances.

They complement their firm data with information on the patents held by these firms, which they identify using the NBER patent data project. The NBER patent data project has information on the primary technology class for all US patents granted from 1963 to 1999 (updated to 2006 since the paper was released), and the number of citations a patent receives, for patents granted in 1975-1999. Most importantly, the dataset provides identifiers to link patents to firms in Compustat.

After excluding firms with no patents or less than four observations over 1980-2001, they have a set of 715 firms, patenting in 426 different technology classes and active in 597 different industries. The typical firm has sales in 5.2 different industries and obtains 16.2 patents per year (mean), though with very high variability. The firms are big, employment of 3,839 people for the median year and firm.

Methodology

Construction of the main explanatory variables of interest has already been described in the empirical contributions section and so here we instead focus on the regressions. There are three regressions of interest described in this digest.

First, their patent regression has citation-weighted patents applied for in each year by each firm as the dependent variable. The main explanatory variables of interest are the log of the previous year's *SPILLTECH* and *SPILLSIC*. In some specifications they also include as additional controls the log of the previous periods R&D stock, log of the previous periods patents, and firm and year specific fixed effects. Because patents and the citations they receive are integers, they use a negative binomial regression (a functional form designed to handle integer explanatory variables).

Second, for their productivity regression, their dependent variable is the log of their measure of quality-adjusted output (sales divided by a price index). They perform an ordinary-least squares regression, with the main variables of interest being the log of the previous year's *SPILLTECH* and *SPILLSIC*. They also include the log of the previous period's R&D stock, the log of their measure of capital, and the log of employment (their measure of labor). As additional controls, they include measures of industry wide output, which should help account for transitory industry-specific shocks. They typically include a set of firm and year fixed effects as well.

Third, for their market value regression, the authors follow a standard procedure for estimating the determinants of market value. The dependent variable is the log of market value divided by the stock of non-R&D assets; that is, the dependent variable can be interpreted as the percent markup over the value of the firm's non-R&D assets. This is also known as "Tobin's average Q" in the literature. The two key explanatory variables are the log of the previous year's *SPILLTECH* and *SPILLSIC*. However, following the procedure that is standard in this literature, an additional explanatory variable is a non-linear function of the R&D stock divided by non-R&D assets. Lastly, as additional controls they sometimes include firm and year fixed effects, as well as the dependent variable lagged by one year.

Lastly, for their regressions that leverage changes in tax policy to identify plausibly random variation in R&D spending, they construct firm and year specific R&D tax rates. As noted above, these rates will vary by firm, as well as over time as tax policy changes. They then run a regression with firm R&D as the dependent variable. The tax policy, plus firm and time fixed effects are the explanatory variables. Note that since these regressions include a firm fixed effect, changes in predicted R&D across firms are exclusively driven by changes in taxes. They then use these predicted R&D values to generate new measures of *SPILLTECH* and *SPILLSIC*, where changes in spillovers are plausibly determined only by changes in tax policy, not factors possibly correlated with the dependent variable. They use these "instruments" for *SPILLTECH* and *SPILLSIC* in their main regressions.