ARTICLE IN PRESS

Research Policy xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Research Policy

journal homepage: www.elsevier.com/locate/respol

Research paper

Inventing by combining pre-existing technologies: Patent evidence on learning and fishing out

Matthew S. Clancy

USDA Economic Research Service, Washington DC, United States

ARTICLE INFO

JEL classification:
O31
O34

Keywords:
Innovation
Patents
Combinatorial growth
Spillovers
R&D

ABSTRACT

I develop a model of innovation where new technologies are combinations of pre-existing technological components. The model captures two opposing forces. The best ideas are used up (knowledge is exhaustible). However, as firms learn which technologies can be combined, new ideas become feasible (knowledge accumulates). I test the model with more than 80 years of US patent data. Technological components are proxied by 13,517 patent office technology classifications. These are reused and recycled in 10,000 distinct three-component sets. Consistent with a learning/fishing-out dynamic, I show patenting in one set of components is correlated with a subsequent increase in similar patents (sharing two of three components), but a subsequent decrease in identical patents (sharing all three components). I use patent renewal data to show my results are not driven by changes in demand for various technology bundles. My results suggest the positive impact of learning on subsequent patenting is larger than the negative impact of fishing out.

Opposite forecasts for the outlook of innovation currently coexist. In one view, rapid innovation lies ahead: artificial intelligence will reshape the economy (Brynjolfsson and McAfee, 2014; Bostrom, 2014) as we take to other planets (Vance, 2015) and use genetic engineering to control our evolution (Doudna and Sternberg, 2017). But in another view, continuous innovation is an exception, and stagnation is the rule. We have already discovered all the good ideas and, as a consequence, innovation is likely to slow and stall (Cowen, 2011; Gordon, 2016). These views differ in their evaluations of two opposing dynamics in innovation. The first emphasizes innovation as a primarily cumulative process: as we learn more, the applications worth exploring multiply. In this paper, I refer to this as the learning effect. The second view emphasizes that knowledge is more like a finite natural resource extracted by research. This is frequently referred to as the fishing out effect. The outlook for innovation depends on which of these features dominates. Are we fishing out the stock of ideas faster than learning "restocks" it? This is an empirical question and this paper develops a novel methodology to answer it.

Psychologist of creativity Keith Sawyer writes creativity is "a new mental combination that is expressed in the world" (Sawyer, 2012 pg. 7). My starting point is a model of innovation wherein ideas are new combinations of pre-existing technological components. Consider the internal combustion engine as a representative example. While it is a single idea, it can also be viewed as a combination of constituent components: pistons, crankshafts, flywheels, and so on. Each of these components existed prior to the engine, and the engine's discovery required assembling pre-existing constituent components into a combination not previously known (Dartnell, 2014, pg. 201).

This way of thinking about discovery has a long history. Mathematician Henri Poincaré argued, "[T]o create consists precisely in not making useless combinations and in making those which are useful and which are only a small minority." (Poincaré, 1913, pg. 386). Abbott Payson Usher's A History of Mechanical Inventions noted, "Invention finds its distinctive feature in the constructive assimilation of preexisting elements into new syntheses, new patterns, or new configurations of behavior" (Usher, 1929, pg. 11). Schumpeter described the essence of enterprises and entrepreneurship to be "the carrying out of new combinations" (recounted in Weitzman, 1998, pg. 335). This perspective has also been articulated in formal models by Weitzman (1998), Olsson and Frey (2002), Simonton (2004), Olsson (2005), Feinstein (2011), Ghiglino (2012), and Akcigit et al. (2013).

A straightforward interpretation of "fishing out" follows from combinatorial models of innovation. This paper assumes a given combination has a fixed number of distinct applications (i.e., there are only so many ways to combine pistons, crankshafts, flywheels, and so on to obtain something novel and useful), so that the stock of ideas is finite and R&D draws it down.[1] Combinatorial models can also model the cumulative nature of knowledge. My model is most closely related to the concept of "clumps" in Arthur (2009), in which some components (such as pistons and crankshafts) are understood to go together

E-mail address: matthew.clancy@ers.usda.gov.

[1] Weitzman (1998) and Akcigit, Kerr, and Nicholas each incorporate a variant of fishing out effect.

https://doi.org/10.1016/j.respol.2017.10.015
Received 18 November 2016; Received in revised form 24 October 2017; Accepted 31 October 2017
0048-7333/ Published by Elsevier B.V.

Please cite this article as: Clancy, M.S., Research Policy (2017), http://dx.doi.org/10.1016/j.respol.2017.10.015

naturally because they "repeatedly form subparts of useful combinations" (Arthur 2009, pg. 70).[2]

To briefly illustrate the thrust of this paper, consider three technological components, *x*, *y*, and *z*, that may be combined into a new idea *xyz* with some probability and at some cost. The combination is more likely to succeed if researchers can observe prior instances where the components have been combined usefully. This knowledge is modeled by how many times each of the *pairs* of components (*xy*, *xz*, and *yz*) have been combined successfully. However, the exact combination of components *xyz* can only be "discovered" a finite number of times.

In this model, every new idea affects future innovation through both learning and fishing out effects. Suppose a fourth technological component, *w*, is also available. If *xyz* is a successful combination, researchers observe an instance of the pairs *xy*, *xz*, and *yz* being combined. This increases the probability that combinations such as *wxy*, *wxz*, and *wyz* will also succeed, as these combinations make use of the same pairs. This is the positive learning effect, where every discovery makes similar research more attractive. At the same time, one instance of the precise combination *xyz* has been used up by its discovery. This is the negative fishing out effect.

There is a long line of empirical papers in this literature. Ideas are usually proxied by academic papers or patents, and the citations they make to antecedents in different fields determine the extent of recombination in an idea. A few papers (e.g., Fleming, 2001 and Akcigit et al., 2013) instead use the technological classifications directly assigned to patents as proxies for technological components. This is the approach I take.

Much of this literature has looked for correlations between the combinatorial properties of patents/papers and their subsequent citations. Because citations can be interpreted as proxies for knowledge flows, this line of literature can also be interpreted as providing some evidence on the learning effect. Patents/papers that receive more forward citations are ideas from which many future researchers learned something important. A long stream of studies[3] has generally found that recombination is associated with more citations, and therefore (perhaps) learning. Conversely, the extent to which familiar combinations do *not* generate new citations could be read as evidence that these technological domains are fished out. Fleming (2001) provides more direct evidence on fishing out by showing a patent is less likely to be highly cited if its exact combination of subclasses has been patented more often.

This paper differs from the above in several respects. My unit of observation is a specific combination in a particular year, not a paper or patent. My dependent variable is not citations, but the number of patent applications in a given year with a particular combination of technological components, including years in which no patent applications for a given combination are filed (the majority of cases). By looking at the factors correlated with the number of applications with a particular combination, I can measure the empirical import of various variables associated with the combination.

Moreover, my proxies for learning and fishing out allow me to identify these effects separately. I assume a combination is fished out by identical combinations. For example, the combination *xyz* is only fished out by patents assigned the exact set *xyz* (this is also how Fleming, 2001 measures fishing out). However, the learning effect is driven by the number of patents using various pairs of elements in a set (i.e., patents containing any of *xy*, *xz*, and *yz*). This gives me differential variation in learning and fishing out, which I use to estimate their relative magnitudes. The chief contribution of the paper is demonstrating that the learning effect exceeds the fishing out effect.

However, this exercise is only useful to the extent that the underlying model and causal interpretations are correct. My second contribution is providing novel evidence to support the model. I derive and find empirical support for five hypotheses suggested by this paper's model of combinatorial innovation. Additionally, I use patent renewal data to rule out an alternative interpretation of the data, that my measure of "learning" is merely proxying for lagged changes in demand for different technologies.

Third, my use of technology subclasses improves on earlier work by aggregating up to the mainline class. This ensures that technology classifications are non-nested, exhaustive, and comparable. Aharonson and Schilling (2016) have recently explored a similar approach as applied to maps of the technological landscape.

The layout of this paper is as follows. In Section 1, I set up my model and supply four of the five hypotheses that will be tested. In Section 2, I describe the historic US patent data I use in the paper's empirical application. Section 3 describes my econometric methodology. Section 4 presents my results, and evaluates how well they support the four hypotheses developed in Section 1. Section 5 extends the analysis by introducing patent renewal data to both test a fifth hypothesis and to eliminate the alternative hypothesis that my results are driven by demand-side factors. Section 6 compares the size of the learning and fishing out effects. Section 7 concludes with a summary of the paper's contributions and some suggestions for future research.

1. A model of combinatorial innovation

1.1. Model

This section presents a three-step model of R&D and patenting. The first part is a combinatorial model of the R&D process. The second is a simple model of how firms decide which R&D projects to initiate. The third part combines the first two to derive predictions about which ideas are patented.

We begin with a model of the R&D process. There exists a set Q of pre-existing technological components. Ideas are subsets of Q with at least two components. Let a subset be denoted by i. The purpose of R&D is to determine if a set of components results in a viable invention, where "viable" means simply that the invention works, in the sense of meeting the desired technical specifications.

The viability of an invention is a function of how its constituent components interact with each other. I define a scalar measure called *affinity* that measures the state of knowledge about how two components can be usefully combined. Let the affinity at time t between a pair j of components be denoted A_{jt}. There is an unobserved "true" affinity that is time-invariant and measures the true utility of combining technological components; A_{jt} converges to this "true" value as information about how components may or may not be combined accumulates. In particular, A_{jt} increases with the number of examples of viable ideas using component-pair j and decreases with the number of examples of unviable ideas using component-pair j.

Inventions are more likely to be viable (from the perspective of researchers) if their components have high affinity for each other. Let $\overrightarrow{A_{it}}$ denote the vector of affinities of all pairs of set i's components at time t (there will be $n(n\text{-}1)/2$ pairs of components if the set i has n components). The probability that an idea with set i is viable is a function of the affinities between its components:

$$\Phi(\overrightarrow{A_{it}}) \equiv \Pr(i \text{ viable}) \quad (1)$$

where I assume:

[2] Others have modeled the knowledge associated with combinatorial innovation as arising from the discovery of new components (Weitzman, 1998; Akcigit et al., 2013), or the discovery of new combinations that can be repeated to diminishing effect (Akcigit et al., 2013), or the discovery of combinations or ideas that bridge distant technological spaces (Olsson, 2005; Feinstein 2011) and reveal the quality of "nearby" ideas (Jovanovic and Rob, 1990; Kauffman et al., 2000; Auerswald et al., 2000).

[3] See Hall et al. (2001), Fleming (2001), Schoenmakers and Duysters (2011), Schilling and Green (2011), Nemet (2012), Akcigit et al. (2013), Kaplan and Vakili (2015). Nemet and Johnson (2012) is an example of a contrary finding. Uzzi et al. (2013) and Keijl et al. (2016) suggest it is not necessarily the total amount of recombination that matters, but that an atypical combination was made within a familiar context.