The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?"

by Benjamin Jones

in *The Review of Economic Studies* 76(1): 283-317. (2009)

Research Digest by Matthew Clancy

Theoretical Contribution

Jones embeds a new model of innovation into an existing model of endogenous growth. New productive ideas (i.e., technologies used in production) are built out of foundational knowledge, acquired in education. As a modeling strategy, required foundational knowledge is represented with the metaphor of a cylinder. All cylinders have the same circumference and arcs around the cylinder correspond to different forms of required knowledge. Longer arcs correspond to a wider "breadth" of knowledge. Figure 1 illustrates the idea. Suppose making a new airplane design requires knowledge of fluid mechanics, thermodynamics, and material sciences, as illustrated. Notice that each field of foundational knowledge "wraps" partway over the cylinder.

Cylinders can vary in their height and technologies that require "more" foundational knowledge have taller cylinders. The height of a cylinder represents the "depth" of required knowledge to innovate. Figure 1 also illustrates this. While developing airplanes always requires knowledge of thermodynamics, fluid mechanics and material sciences, during the era of the Wright brothers, we simply didn't know much about these fields and it was possible invent airplanes without much foundational knowledge (short cylinder in Figure 1: left). Since then we have learned a ton about each area, and new airplane designs fully build on this larger store of foundational knowledge. Modern airplane design requires access to much more knowledge (tall cylinder in Figure 1: right).

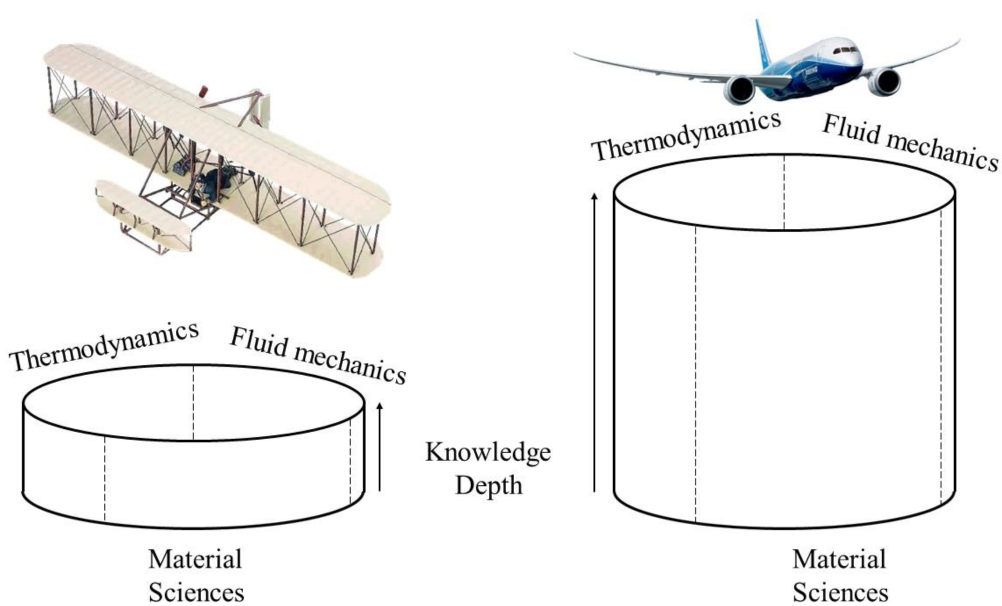

Figure 1. Depth and Breadth of Knowledge

The cost of acquiring knowledge relevant for innovation in technology field j (represented by E_j) is represented by the following function that multiplies breadth (b_j) and depth ($D_j(t)$) and raises the product to a (positive) exponent ε:

$$E_j = (b_j D_j(t))^{\varepsilon} \tag{0.1}$$

Note the depth required can vary with time t, and that the cost of education is rising in both breadth and depth.

Innovators decide how much breadth of knowledge they will acquire, represented as how much of today's knowledge cylinder they will wrap around. For example, in Figure 1 an innovator could decide to acquire knowledge wrapping 100% of the way around the cylinder, thus enabling them to invent new airplanes on their own. Or they could decide to obtain knowledge wrapping 1/3 of the way around the cylinder, specializing in something like "Fluid Mechanics." In the latter case, they will need to team up with other innovators whose knowledge complements their own, so that collectively their knowledge spans the required breadth, if they want to invent new airplanes.

Once an individual or team has enough foundational knowledge to innovate, the model describes how they may invent and patent new technologies. New technologies are licensed to producers, and innovators are compensated with licensing fees. Both the size of the economy and the depth of knowledge are increasing in the number of technologies available. The model is completed with a description of how each worker decides (1) whether to be a production worker or innovator and (2) if an innovator, what field to specialize in and what breadth of knowledge to obtain (they always go as deep as necessary to invent new things).

Along a balanced growth path, a series of equalities lead to the paper's main results. Under balanced growth, the share of the population engaged in innovation is neither growing nor shrinking. This means the expected value of being a worker must equal the expected value of being an innovator in any field j; otherwise, people would switch from one occupation to the other. Let V represent the lifetime value of a profession:

$$V(\text{worker}) = V(\text{field } j \text{ innovator}) = V(\text{field } i \text{ innovator}) \tag{0.2}$$

The same principle implies the value of being an innovator is the same in every field. The expected value of being an innovator is the net value of expected innovation licensing fees less the cost of education. Crucially, Jones assumes that innovation is *isoelastic* in breadth of knowledge, meaning a 1% increase in breadth leads to a fixed percentage increase in new ideas (whether your breadth stretches ¼, ½, or whatever around the cylinder). He shows that when this assumption is true, the cost of education will always be equal to a constant proportion of lifetime value. Since every field has the same expected value, this implies workers spend the same amount on education *for every field*. Let s be the share of lifetime income spent on education. Then:

$$E_j = sV(\text{field } j \text{ innovator}) = sV(\text{field } i \text{ innovator}) = E_i \qquad (0.3)$$

But different fields may have different knowledge depths. Some fields may be represented by short cylinders, like in Figure 1: left, and others by tall cylinders, like in Figure 1: right. If everyone spends the same amount on education though, inventors in fields with very deep knowledge requirements must have very narrow expertise, and inventors in fields with very shallow knowledge requirements must have very broad expertise. This can be seen if we rearranging equation (0.1) and observe that $E_j = E$ for any field j:

$$b_j = \frac{1}{D_j(t)} E^{1/\varepsilon} \qquad (0.4)$$

We can summarize the main theoretical results as follows.

1. At any point in time, innovators all spend the same amount on education, as indicated by equation (0.3).

2. At any point in time, innovators in fields with deeper knowledge requirements obtain narrower breadth of their knowledge (they specialize more), as indicated by equation (0.4).

3. Because workers in fields with deeper knowledge requirements are more specialized (have less breadth), bigger teams are required to innovate in these fields.

4. Because workers spend a constant share of their lifetime income on education, as the economy grows in per-capita terms, the costs of education rise.

5. If the depth of knowledge grows faster than the economy, workers specialize more and more over time.

6. Because workers specialize more and more over time, the size of teams grows over time.

Jones concludes with an explicit derivation of the economy's growth rate. This is proportional to the growth of the labor supply. Growth is faster the more the productivity of research improves in response to new knowledge. Growth is slower the faster the depth of knowledge required for innovation grows.

Empirical Findings

Jones uses data on US patents granted between 1975 and 2001 to complement his theoretical contributions. He develops several proxies for team size, cost of education, specialization, and knowledge depth.

The simplest proxy is team size, for which Jones simply uses the number of inventors listed on each patent. He shows the average number of inventors per patent has risen from 1.7 to 2.3 between 1975 and 1999, consistent with theoretical result #6 above.

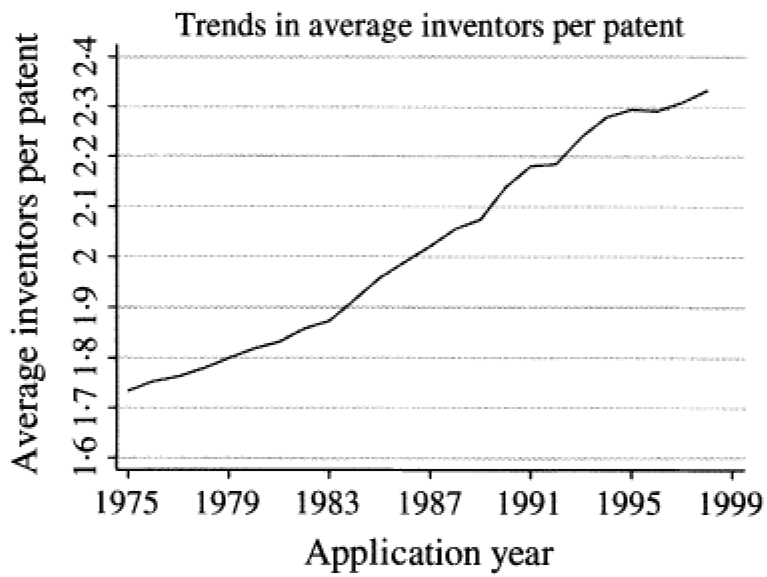

Figure 2. Inventors per patent on the rise

The next simplest proxy is for the cost of education. Because education and training takes time, Jones uses the age of first patent for the cost of education. As it becomes necessary to go to college, and then to obtain a masters, and then a PhD before one knows enough to invent, the age of first invention should rise. Jones documents the age of first invention has risen from about 30.5 to 31.5 between 1985 and 1998, consistent with theoretical result #4 above.

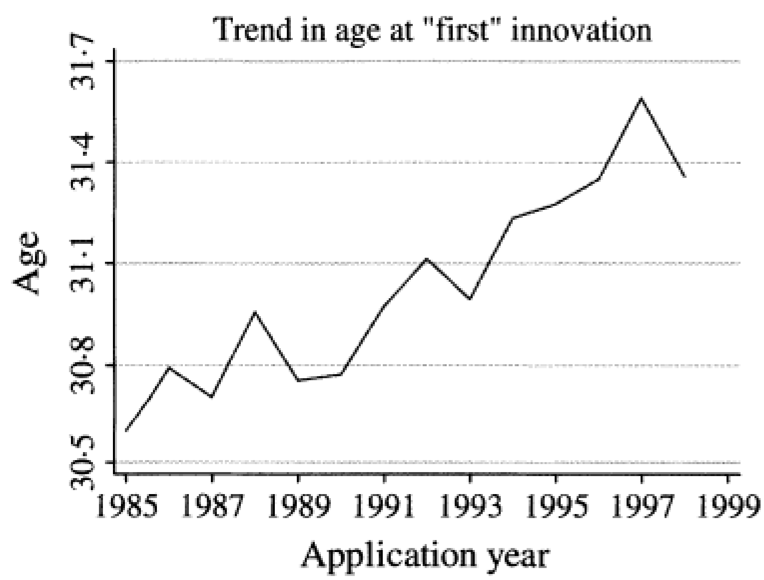

Figure 3. Age at First Invention Rising

Jones' proxy for specialization relies on the US Patent and Trademark Office's patent classification scheme. Patents are assigned to one of 414 major technological categories, and

Jones reasons that it is harder for very specialized workers to work in different categories. The probability of patenting in differing technology categories between subsequent inventions is his measure of specialization. Jones shows the probability of switching technology categories between consecutive patents (granted within 3 years of each other) has declined from 0.57 to 0.51 over 1975 to 1993, consistent with theoretical result #5 above.

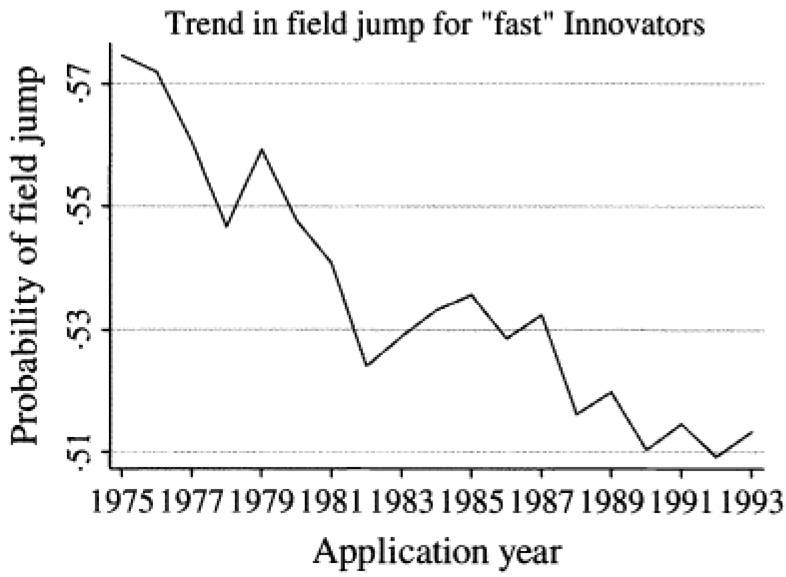

Figure 4. Probability of Switching Fields Falling

Finally, Jones' construction of a measure of "knowledge depth" is considerably more complicated than the others and I reserve a full discussion of it to the methodology section. But in brief, patents make citations to other patents. If these citations are interpreted as a signal that a patent is using ideas from the cited patent, then the more citations a patent makes (and the more citations the *cited* patents make), then plausibly the more knowledge went into the technology.

Jones runs a number of regressions linking knowledge depth to his other variables of interest. Knowledge depth is positively correlated with the number of inventors on a patent, consistent with theoretical result #3 above. Knowledge depth is negatively correlated with the probability of switching technology categories, consistent with theoretical result #2 above. And knowledge depth is uncorrelated with the age of first invention, consistent with theoretical result #1 above.

Last, as a further check, Jones compares the number of inventors per patent and age at first invention across different technology fields. As he predicts, there is little variation in the age of first invention (consistent with theoretical result #1 above), but there is significant variation in the number of inventors per patent (consistent with theoretical result #3 above).

Discussion

The presentation of this digest is a bit misleading. In the paper, Jones begins with his empirical findings and then develops a theory that ties them together. His develops a model where the "burden" of knowledge necessary to innovate continually grows over time. Consistent with his patent data, inventors respond by specializing more, relying on teams more, and spending more time receiving an education. Thus there is a growing "burden of knowledge" and the decline and death of "renaissance men" who are capable of making significant contributions in many different fields.

Data

Jones' empirical data is drawn from the NBER patent data project (https://www.nber.org/patents/), from Jaffe, Hall, and Trajtenberg. This includes data on all 2.9mn US patents granted between 1963 and 2001, with more detailed data on the 2.1mn patents granted since 1975.

He supplements this data with data on birthdays from the website www.AnyBirthday.com. At the time of writing, the website had data on 135 million Americans' birthdays, gleaned through public records. Jones uses names and zip codes to match inventors on patents to individuals on the website and obtains 56,281 unique matches (one inventor linked to a single name). Notably, the sample of inventors whose birthdays are matched is not a random sample of the population of inventors, but is skewed towards inventors who did not assign their patent to another organization (or individual) at the time of grant. This is because inventors who assign their patents to an organization (generally their employer) do not provide their zip codes as often as "independent" inventors. Nonetheless, restricting attention to the patents of inventors whose birthdays Jones observes reveals broadly similar trends.

Methodology

There are a few subtleties in the construction of Jones' proxies. When constructing age at first invention, Jones cannot observe patents granted before 1975, which calls into question whether observed patents are really the first invention of the listed inventor. To address this, he restricts his analysis to inventors between the ages of 25 and 35 whose first observed patent comes after 1985.

Second, when constructing his measure of specialization, Jones looks for field switches that occur between consecutive patents applied for within 3 years of each other. However, we only observe applications that were granted before the dataset ends in 1999, and there can be considerable variation in the time required to move a patent from application to grant (usually 2 years, but possible much longer). Jones therefore limits his observations to applications that were granted in 3 years or less.

Third, to construct a proxy for knowledge depth, Jones relies on patent citations. Interpreting these citations as signals of knowledge flows (for example, Patent X might cite patent Y if it uses some of the ideas associated with patent Y), Jones uses the size of the "citation tree" as a proxy for knowledge depth. The citation tree is the total number of citations a patent makes, *plus* the citations made by the patents it cites, *plus* the citations made by those patents, as far back as observations permit. Patents with a large citation tree tend to cite lots of patents that, in turn, cite lots of patents.

There are two challenges with the raw citation tree. First, the trees vary tremendously in size. To avoid letting results be dominated by a small number of very large citation trees, Jones uses the log of tree size. Second, tree size grows over time as more data becomes available and this does not necessarily represent greater knowledge depth. To adjust for this, Jones computes the mean value of the tree size for each year, and he measures relative knowledge depth as the difference from this mean value, measured in standard deviations for that year. For example, suppose a patent granted in 1995 has a citation tree with 2,981 citations. The natural log of this is 8. Suppose in 1995 the mean log citation tree is 4 with a standard deviation of 2. This means our patent with a citation tree of 2,981 citations is 2 standard deviations above the mean for 1995. Jones would use "2" as his proxy for knowledge depth in this field in this year.

Jones uses his constructed data to document trends over time and in cross-section. For trends over time, Jones supplements the above figures with OLS regressions (in the case of team size and age at first invention) or probit regressions (in the case of probability of switching fields). These allow him to adjust for possibly confounding changes in technology field and the rise of foreign patenting. In all cases, the trends documented in his diagrams remain robust. He also uses OLS and probit regressions to investigate the correlation between knowledge depth and his other proxies as described in the empirical findings section. He demonstrates his results are robust to additional control variables such as inclusion of the squared knowledge depth measure, and the number of direct citations make.