“Artificial Intelligence in Research and Development" by Benjamin F. Jones

Discussion by Bronwyn H. Hall

6 September 2025

Jones’ paper is addressed to an interesting question: how will the use of AI affect the productivity of research and development? He addresses this question by developing a task-based model of the R&D process. This model posits that the impact of artificial intelligence on research and development is determined by three principal factors: the proportion of research tasks that AI can undertake, the efficiency of AI in performing these tasks, and the degree of "bottlenecks" inherent in the innovation process. To explore these factors, Jones presents a model that maps R&D inputs to progress in a specific outcome. This model conceptualizes R&D as a series of tasks that can be performed by either humans or machines and derives the implications for growth when machines become more productive via AI or via an increase in the number of tasks they can perform.

This discussion has three sections: 1) Some puzzles with the model, which appears to contain an assumption that renders independent analysis of changes in machine productivity and the spread of machines across tasks problematic. 2) In the computations of potential impacts of AI, a capital share is used that is considerably at variance with the evidence from many countries and time periods. 3) An historical example that might be useful: the changes induced by the enormous fall in computing costs on research in particle physics over the past seventy years.

The model

The model is designed to separately analyze the impact of increases in the share of tasks that machines perform (γ_t) and the impact of increases in their productivity at these tasks (M_t). Unfortunately, it is not really possible to separate these two, as the share of tasks that machines perform will be affected by changes in their productivity at tasks. So I was somewhat puzzled by parts of the analysis.

M_t is an index of machine productivity, formed as a CES of individual machine productivities with coefficient (θ-1)/θ, where θ measures the degree of complementarity in the tasks:

$$M_t = \left[\frac{1}{\gamma_t} \int_0^{\gamma_t} m_t(j)^{\theta/(1-\theta)} dj \right]^{\frac{1-\theta}{\theta}}$$

The share of tasks that will actually be performed by machines is given by the technology adoption condition, which specifies that the productivity per cost of the machine in the task must be greater than labor productivity in the task relative to the wage. Assumption 1 in

the paper makes it clear that γ_t is not the share of tasks that might be performed by machines, but is the share that actually will be performed by machines:

$$m_t(j) \geq H\mu_t / w_t \text{ for all } j \in [0, \gamma_t]$$

That is, if we rank the $m_t(j)$ by their value from high to low, γ_t is the value of j for which $m_t(j) \geq H\mu_t / w_t$ and $m_t(j+1) < H\mu_t / w_t$. The implication is that M_t and γ_t are jointly determined from the primitives, which are the set of productivities m_ts and their cost μ_t (in addition to the labor productivity and cost).

Does this matter? For the static analysis of small short run changes, it seems fairly innocuous to examine the impact of an increase in "average" machine productivity holding constant the share of tasks performed by machines. But for the larger changes considered in section 4.1, it seems a bit problematic. For example, Corollary 3 examines the impact of a large jump in M_t, which must imply that at least some of the m_ts experience large increases. This in turn would increase the share of tasks that machines perform, so that treating γ as a parameter, as the Appendix proof does, might be inappropriate.

Reconciling the model with the data

The discussion below will focus on the problem of bringing the model to the data available. In particular, can we learn anything from the existing data on R&D costs together with Jones' model about the impacts of AI on R&D productivity? I begin with some comments on Jones' choice of a value for the capital cost share in R&D spending (one third), which I found quite high when compared to numbers in the literature and my prior experience with the data. Fortunately there is fairly recent data on the R&D cost breakdown for quite a few countries.

The OECD has data on the cost breakdown of business R&D for a number of developed economies. These data are currently collected under a standard defined by OECD's Frascati Manual (2015), which defines roughly 15 different cost categories, some of which are subcategories of others. Earlier versions of the Frascati Manual and earlier data collections prior to 2015 generally specified only three cost categories: capital, internal labor expense, and other expenditures. These are enough to get a broad picture of the evolution of the total capital share in R&D expenditure. The average of this share over the countries for which data are available is shown in Figure 1. Clearly the average capital share is well below one third, and in fact is declining in the recent past. Note that this share includes capitalized software, which is where we expect some of the AI investment will be counted, at least in the recent past.[1]

[1] Accounting for R&D-related capitalized software has been introduced into national income accounting at different times in different countries, although the Frascati Manual makes it clear that it should be included in R&D capital, at least since the 2002 edition.

Figure 1

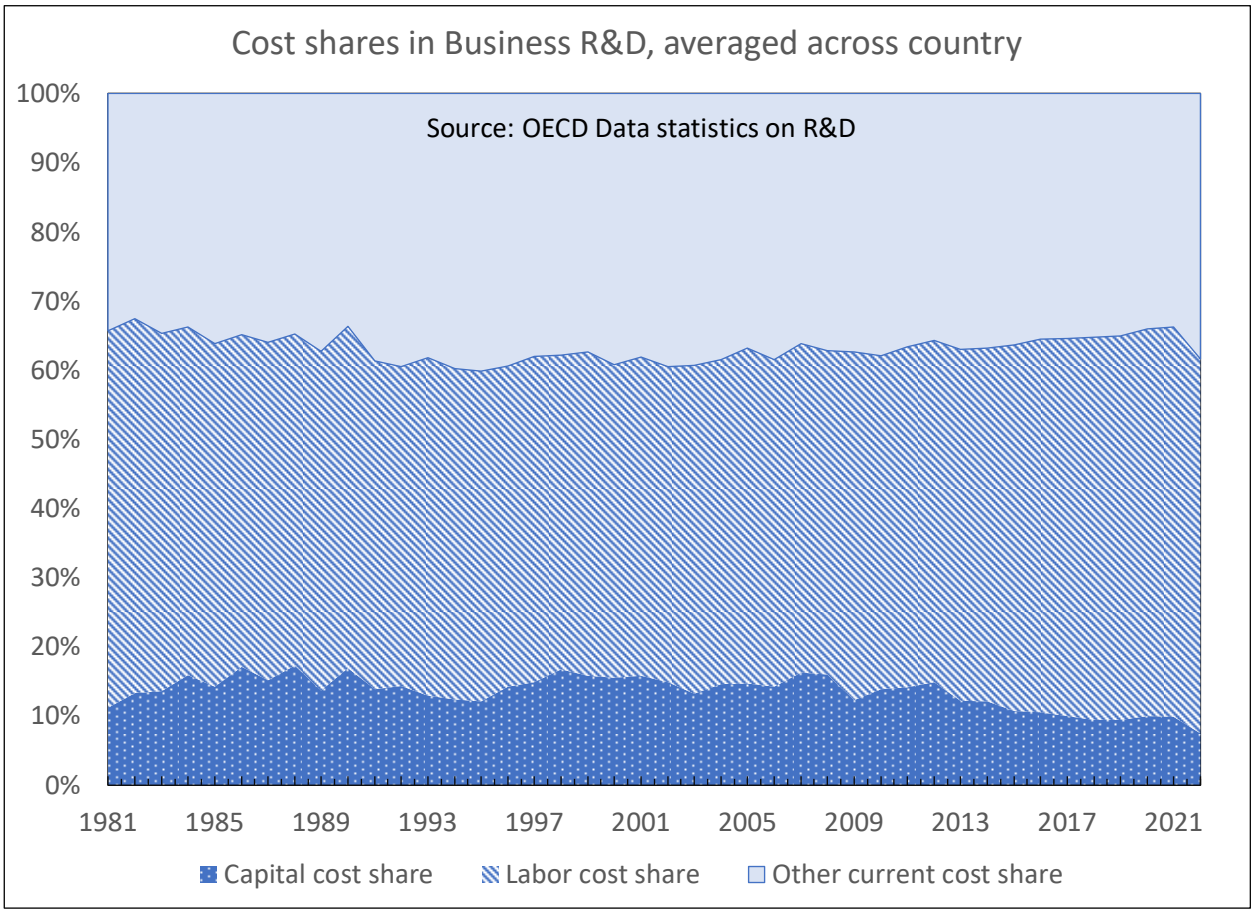

Unfortunately, the United States is a latecomer to the detailed collection of the components of R&D spending, only entering the data in Figure 1 in the years after 2018. I show more detail on the cost breakdown for R&D in the US using data from the latest year available, 2021, in Table 1. These figures show an even lower share of capital in R&D cost.[2] The conclusion is that using existing data, a capital share of about 10 per cent would be more appropriate than one third; in fact, if we exclude land and buildings, it may even be an overestimate.

[2] The capital cost share is even lower for R&D funded by the government, hovering around 1-3 per cent of costs over the whole period. It is also possible that some capital cost is included in the measures of current cost, for example the rental of capital equipment, but it seems unlikely that this will be enough to increase the capital share to one-third.

Table 1

Composition of Business R&D spending in the US - 2021	
	Share
Total capital cost, of which	**8.25%**
Other IP products (purchased)	0.89%
Land & buildings	0.91%
Machinery	3.07%
Software	2.54%
Other	0.84%
Total current cost, of which	**91.75%**
Labor costs	61.71%
External R&D personnel	4.20%
Current cost NEC	15.94%
Materials purchased	6.53%
Services purchased	3.36%
Total	**100.00%**

Using 10 per cent as the capital share, we can recompute some of the numbers in Section 4. For example, an infinite increase in machine intelligence would increase the rate of progress by a factor of 1.23, not 2.25. If only half of the machines have an infinite increase in intelligence, the increase in the rate of progress would be 1.11, not 1.44.

On the other hand, the impact of a complete shift to AI machines on the rate of progress would be 25 rather than 2.25. This is because the capital share is 0.1 but it is assumed that the initial potential share of machine tasks is 0.5. With the bottleneck parameter of 2, we have $(0.5/0.1)^2 = 25$. This happens because machines are assumed to be doing half the tasks initially, but only account for one-tenth of the expense, so they are very productive relative to labor. If all the tasks are shifted to machines, there is a huge boost in the rate of progress.

An example

In order to provide a detailed example of the impact of automation on research, this section describes the evolution of fundamental research in an area with which I am somewhat familiar: "bump hunting" in high energy physics. This term describes the practice of looking for new particles or new phenomena by repeating the same interaction many times using some kind of particle accelerator and then examining the distribution of outcomes statistically. Beginning in the 1960s as an essentially manual procedure once the relevant experiment had been conducted, this type of research has evolved to use machine learning

and AI to become (almost) fully automatic. Much of the improvement in efficiency has been due to the fall in computing cost and the shift away from humans to software for analysis.

In the 1960s and early 1970s, bump-hunting was a fundamentally human-centric activity. Physicists would meticulously plot experimental data by hand or with early computer programs, creating histograms of particle interaction properties such as invariant mass. They would then visually scan these plots for interesting deviation from what they expected to see (normally a rather smooth distribution) and identify these with new or known particles. The first step in this process was to define an interesting experiment that would be repeated thousands of times so that a distribution of the possible results could be collected. For example, a typical experiment at the Lawrence Radiation Laboratory (now the Lawrence Berkeley Laboratory) Bevatron involved firing a proton into the bubble chamber filled with liquid hydrogen and photographing its track as it interacted with the bubbles, possibly producing other particles that then decayed.

The experimental data itself was often collected by using the FSD (Flying Spot Digitizer) on each photograph of an interaction (of which there might be thousands). The operators responsible for this were skilled at the task, but were not trained physicists, nor were they required to even have a bachelors degree. The resulting data were then fitted by a software program to yield information on the particles emitted and their energies. Distributions of the results were plotted, usually manually, and examined for evidence of anomalies that would indicate the presence of a new particle or phenomenon. For some examples of this kind of research, see Bland et al. (1968) and Goldhaber et al. (1969).

Over time, several advances were made in this process: first, statistical techniques were developed to deal with the problem of false positives and make the finding and testing for anomalies more rigorous.[3] Second, the bubble chamber was replaced with other kinds of detectors that could operate at much faster rates and provided electronic readout, obviating the need for FSD operators. For example, the Nobel prize-winning discovery of the charm quark J/ψ simultaneously at the Stanford Linear Accelerator (SLAC) and Brookhaven National Laboratory was done by the SLAC team using a series of spark chambers that recorded the results of the interactions of positive and negative electrons. (Augustin et al. 1974, 1975). The Brookhaven team used a proton-proton interaction and spectrometers to find the same result (Aubert et al. 1974).

Third and more important for the shift towards the use of machines and software, advances in computing allowed a number of tasks formerly done by humans to shift to machines and for new methods to become possible. For example, computational speed improvements allowed rapid scanning of vast amounts of data in search of rare "resonances," that is, particles not yet discovered. For example, the Collider Detector at Fermilab (CDF) was used

[3] The problem should be familiar to economists – as the number of signal searches increased, the number of false signals also increased at any fixed significance level. The solution used in high-energy physics is to simulate the background distribution from known physics many times and compute how often a random significant bump would arise.

to analyze approximately 10^{14} proton-antiproton interactions by searching for statistical anomalies in their outcomes relative to the outcomes predicted by the standard model of particle physics. This exercise produced a negative result as it detected only the number of anomalies predicted by simulation of the standard model (CDF Collaboration 2009).

The current frontier technique in this area augments the task with machine learning, implying that AI methods have already been introduced in this area. The successful search for the Higgs Boson at the CERN Large Hadron Collider made extensive use of machine learning to classify the outcome of the particle collisions and select those for further analysis (ATLAS 2012; CMS 2012). A glance at the more recent publications in this area gives an idea of the increase in complexity of the analysis since the 1960s (Collins et al. 2018, 2019).

How does this historical development of a research history relate to the model in the paper? For example, we might define the desired outcome Z as the discovery of new particles. However, as the perceived difficulty of finding yet another particle increased over time ($\varphi<0$), clearly the budget D was not held constant. Besides the enormous increase in the cost of the necessary particle accelerator over time,[4] the number of authors on the relevant papers cited increased from a handful to several hundred. It is probable that this dynamic (diminishing returns and increased budget) is characteristic of many individual lines of research, as suggested by the examples in Bloom et al. (2020).

A second observation is that even though increased computing speed and other technical improvements meant that much of the original effort devoted to recording the experimental data and analyzing it moved from humans to machines this was accompanied by the expansion of tasks to manage both the experiment and the analysis of the results. That is, the impact of improvements in machine technology was mainly to increase the difficulty of the research undertaken and the cost of that research. in this respect the activity is not unlike other not-for-profit endeavors such as health care and performing arts (Bowen and Baumol 1965).

A last related observation is that although the impact of the increase in computing speed and detector design over the past 70 or so years has meant that many tasks have moved from humans to machines, many new tasks have arisen in making this transition, many of them in the form of software design done by humans. It remains to be seen whether AI can help with this task or will simply make the research more ambitious in both cost and desired outcomes.

Conclusion

[4] Initial costs in 2025 dollars for the Bevatron was $110 million, for Fermilab's Tevatron $375 million, and for CERN's Large Hadron Collider $8.5 billion. These are of course approximate figures and do not include important later enhancements.

This discussion has examined Jones' task-based model of AI's impact on research and development productivity in three ways: theoretical consistency, plausibility of empirical assumptions, and a historical example. While Jones' framework provides a valuable foundation for understanding how AI might transform R&D processes, several important issues warrant attention before the model can effectively inform policy or investment decisions.

The theoretical analysis reveals some tensions in the model's structure that make it difficult to follow the arguments that lead to the key results. In the current version of the paper, the joint determination of machine task share (γ_t) and machine productivity (M_t) creates analytical complications that make it difficult to isolate the effects of different types of AI advancement. I hope that future versions will clarify these issues.

More critically, the empirical calibration of the model appears to significantly overstate the capital intensity of R&D activities. OECD data consistently show capital shares in business R&D of approximately 10 percent, not the one-third assumed in Jones' calculations. This discrepancy has implications for the model's predictions: using the empirically supported capital share reduces the estimated impact of infinite machine intelligence because the capital share is now lower. However, it also suggests that if AI can indeed shift a large proportion of tasks from labor to capital, the productivity gains could be even more dramatic than Jones initially estimated, because capital has been assumed to be much more productive than its current share.

The historical example from high-energy physics illuminates both the promise and complexity of automation in research. The evolution of particle physics from manual data analysis in the 1960s to machine learning-enabled discovery of the Higgs boson demonstrates that technological advance in the form of computing speed and software development did transform research methodologies fundamentally and we might expect similar change from further development of AI methods. However, the past transformation was accompanied by expanding research scope, increasing complexity, and growing team sizes—dynamics that complicate simple productivity calculations. The physics example suggests that AI's impact on R&D may be characterized less by pure efficiency gains and more by the ability to tackle previously intractable problems, a distinction that may require different modeling approaches.

Looking forward, several research priorities emerge from this analysis. First, more detailed empirical work is needed to understand the actual task structure of R&D across different fields and the potential for AI substitution in each category. Second, the model would benefit from incorporating dynamic elements that capture how research frontiers expand as new tools become available—the moving target problem illustrated by the particle physics case. Third, the relationship between AI adoption and research team structure deserves attention, as automation may change not just individual productivity but also the optimal organization of research efforts.

References:

ATLAS Collaboration (2012). Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters* B 716(1): 1-29. https://doi.org/10.1016/j.physletb.2012.08.020

Aubert, J. J., U. Becker, P. J. Biggs, J. Burger, M. Chen, G. Everhart, P. Goldhagen, J. Leong, T. McCorriston, T. G. Rhoades, M. Rohde, S. C. C. Ting, S. L. Wu and Y. Y. Lee (1974). Experimental Observation of a Heavy Particle *J*. *Physical Review Letters* 33(23): 1404-1406. https://doi.org/ 10.1103/PhysRevLett.33.1404

Augustin, J. E., A. M. Boyarski, M. Breidenbach, F. Bulos, J. T. Dakin, G. J. Feldman, G. E. Fischer, D. Fryberger, G. Hanson, B. Jean-Marie, R. R. Larsen, V. Lüth, H. L. Lynch, D. Lyon, C. C. Morehouse, J. M. Paterson, M. L. Perl, B. Richter, P. Rapidis, R. F. Schwitters, W. M. Tanenbaum, F. Vannucci, G. S. Abrams, D. Briggs, W. Chinowsky, C. E. Friedberg, G. Goldhaber, R. J. Hollebeek, J. A. Kadyk, B. Lulu, F. Pierre, G. H. Trilling, J. S. Whitaker, J. Wiss and J. E. Zipse (1974). Discovery of a Narrow Resonance in *e+ e-* Annihilation. *Physical Review Letters* 33(23): 1406-1408. https://doi.org/10.1103/PhysRevLett.33.1406

Augustin, J. E., A. M. Boyarski, M. Breidenbach, F. Bulos, J. T. Dakin, G. J. Feldman, G. E. Fischer, D. Fryberger, G. Hanson, B. Jean-Marie, R. R. Larsen, V. Lüth, H. L. Lynch, D. Lyon, C. C. Morehouse, J. M. Paterson, M. L. Perl, B. Richter, R. F. Schwitters, F. Vannucci, G. S. Abrams, D. Briggs, W. Chinowsky, C. E. Friedberg, G. Goldhaber, R. J. Hollebeek, J. A. Kadyk, G. H. Trilling, J. S. Whitaker and J. E. Zipse (1975). Measurement of $e^+ e^- \rightarrow e^+ e^-$ and $e^+ e^- \rightarrow \mu^+ \mu^-$. *Physical Review Letters* 34(4): 233-236. https://doi.org/10.1103/PhysRevLett.34.233

Baumol, W. J. and W. G. Bowen (1965). On the Performing Arts: The Anatomy of Their Economic Problems. *American Economic Review* 55(1/2): 495-502.

Bland, R. W., G. Goldhaber, B. H. Hall, and G. H. Trilling (1968). Search for A K* Below Kπ Threshold. *Physical Review Letters* 21(3): 173-174. https://doi.org/10.1103/PhysRevLett.21.173

Bloom, N., C. I. Jones, J. Van Reenen and M. Webb (2020). Are Ideas Getting Harder to Find? *American Economic Review* 110(4): 1104-1144. https://doi.org/10.1257/aer.20180338

CDF Collaboration, (2009). Global search for new physics with *2.0 fb^{-1}* at CDF. *Physical Review* D 79(1): 011101. https://doi.org/10.1103/PhysRevD.79.011101

CMS Collaboration(2012). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters* B 716(1): 30-61. https://doi.org/10.1016/j.physletb.2012.08.021

Collins, J. H., K. Howe and B. Nachman (2019). Extending the search for new resonances with machine learning. *Physical Review D* 99(1): 014038. https://doi.org/10.1103/PhysRevD.99.014038

Collins, J., K. Howe and B. Nachman (2018). Anomaly Detection for Resonant New Physics with Machine Learning. *Physical Review Letters* 121(24): 241803. https://doi.org/10.1103/PhysRevLett.121.241803

Goldhaber, G., W. R. Butler, D. G. Coyne, B. H. Hall, J. N. MacNaughton and G. H. Trilling (1969). $\omega\rho^0$ Interference Effect in the $\pi^+p \rightarrow \pi^+ \pi^- \Delta^{++}$ Interaction. *Physical Review Letters* 23 (23): 1351-1354. https://doi.org/10.1103/PhysRevLett.23.1351

OECD (2015). *Frascati Manual 2015: Guidelines for Collecting and Reporting Data on Research and Experimental Development, The Measurement of Scientific, Technological and Innovation Activities.* Paris: OECD Publishing. http://dx.doi.org/10.1787/9789264239012-en

OECD (2025). *Data Explorer*. Science, Technology, and Innovation indicators: R&D statistics. https://data-explorer.oecd.org/