Research Policy 52 (2023) 104665

Contents lists available at ScienceDirect

Research Policy

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

Short communication

Are ideas being fished out?

Leonardo Klüppel [a], Anne Marie Knott [b,*]

[a] Ohio State University, 856 Fisher Hall, 2100 Neil Ave, Columbus, OH 43210, United States of America
[b] Washington University in St Louis, One Brookings Drive, Campus Box 1133, St Louis, MO 63130, United States of America

ARTICLE INFO

JEL classifications:
D24
E23
O30
047

Keywords:
Endogenous growth
Research productivity
Fishing out

ABSTRACT

This paper examines whether declining research productivity can be explained by fishing out—is the production of new knowledge decreasing in the level of existing knowledge? We estimate the knowledge production function for US firms and find instead that knowledge production is increasing in the knowledge stock. This is reinforced by the observations that maximum research productivity across firms is increasing over time, and that research productivity year effects continue to exhibit decline after modeling contributions from knowledge and research labor. Given that fishing out appears unable to explain the decline in research productivity, we offer preliminary evidence of contingent factors that might contribute to the decline.

We have known for several decades that innovation is the primary source of economic growth (Solow, 1957). While Solow was silent on where innovation itself comes from, Romer (1990) articulated a theory in which growth is endogenously driven by purposeful investment in R&D. An important proposition from Romer's theory is a "scale effects prediction" that growth should be proportional to the level of R&D.

This prediction held in the United States from 1950 until 1980. Nominal GDP growth tracked the rise in R&D spending in the 1950s and early 1960s, as well as the decline from the mid-1960s to the late 1970s (Fig. 1). However, the relationship fails to hold after that. As R&D rose in the 1980s, GDP growth never followed. In fact, it has declined ever since.

The leading explanation for the recent failure of the scale effects prediction is that knowledge is subject to "fishing out" (Jones, 1995). The intuition for this theory is that there is a finite number of ideas and, as the best ideas are consumed, we are left with increasingly lower-quality ideas. Technically, this is implemented as decreasing returns to the knowledge stock in the production of new knowledge.

The first implication of Jones's theory, if correct, is that we will always need increasing amounts of R&D to maintain a given level of growth. The second and more important implication is that growth from R&D will ultimately converge toward zero. When it does, no amount of R&D will generate growth—there will be no point in either public or private investment in R&D.

Recently, Jones and colleagues (Bloom et al. (2020) (BJVW) purported to test the fishing out hypothesis. They developed a measure, *research productivity*, calculated as growth in knowledge, divided by the research investment to advance that knowledge. The authors looked within several domains, using a context-specific measure of knowledge for each: semiconductors (transistor density), agriculture (bushels/acre), health care (life expectancy), and pharmaceuticals (new molecular entities), and they found that *research productivity* has declined in each of them.

However, they never tested fishing out directly. Rather they demonstrated that research productivity is declining "virtually everywhere we look" (Bloom et al., 2020, p.1138). In essence, they showed that the macro-level observation of declining research productivity (U.S. GDP growth/U.S. R&D), evident in Fig. 1, holds at micro levels as well.

Because the issue of whether growth from innovation converges toward zero has tremendous implications for both public policy and firm strategy, we attempted to test the fishing out hypothesis directly. We used BJVW's data to test Jones's model structurally. We found no evidence of fishing out. Instead, we found increasing returns to the knowledge stock. This finding resonates with real-world evidence of several important, yet fairly recent, general-purpose innovations like the internet, GPS, and smartphones. We would not expect such important innovations to be so recent if ideas are being fished out.

The finding that knowledge exhibits increasing returns is encouraging, in that it restores the expectation that R&D investment can generate growth in perpetuity. However, it leaves the problem of

* Corresponding author.
E-mail addresses: kluppel.1@osu.edu (L. Klüppel), knott@wustl.edu (A.M. Knott).

https://doi.org/10.1016/j.respol.2022.104665
Received 6 June 2022; Received in revised form 31 October 2022; Accepted 7 November 2022
Available online 24 November 2022
0048-7333/

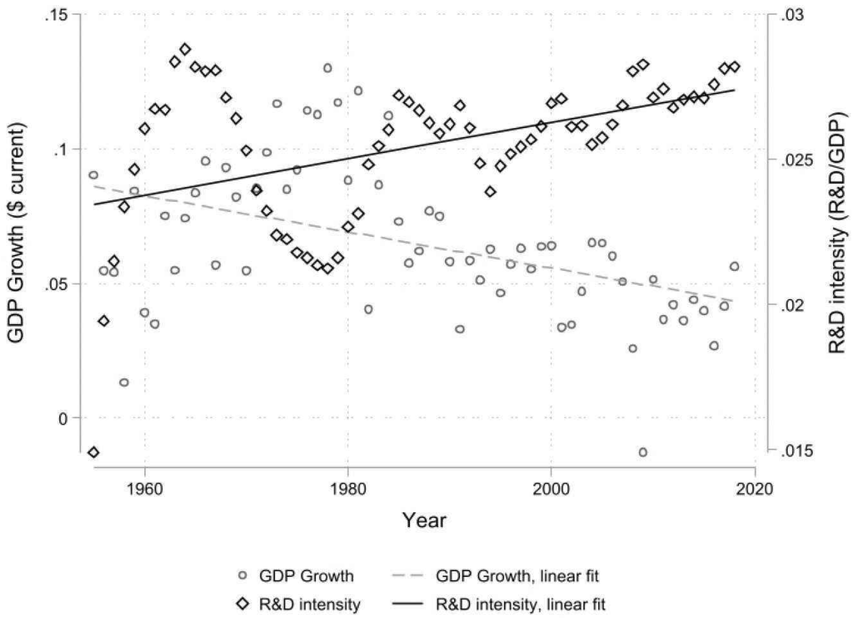

Fig. 1. Relationship between R&D spending and GDP growth

Notes: In this figure we compare U.S. GDP growth (current dollars/current dollars) to U.S. R&D intensity (total R&D from all sources/GDP). GDP growth tracks the rise in R&D spending in the 1950s to the mid-1960s, with a lag. GDP growth also tracks the decline in R&D spending from the mid-1960s to 1980, but it continues to decline even after R&D spending rose in the 1980s. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

declining R&D productivity. We propose that contingent factors are driving the decline. If so, it may be possible to restore R&D productivity and revive growth by addressing those factors. Until then, we will need increasing amounts of R&D merely to maintain even the recent low growth levels we have experienced.

Our work is similar in spirit to Clancy (2018), who built and empirically tested a model of innovation-driven growth that allowed for both “learning” (new combinations of technological components) and “fishing out” (pre-existing combinations of components). He found that learning effects dominate fishing out, which seems consistent with the results here. However, Clancy's definition of fishing out differs from Jones's. Accordingly, his results are not directly comparable to ours.

This work is also similar to that of Furman et al. (2002), who utilized Romer's model to evaluate differences in national patent productivity. Their baseline specification indicated that there are increasing returns to the knowledge stock, which matches our main result. Thus, their results and ours are inconsistent with fishing out.

The paper proceeds as follows. First, we review our conceptual framework, which builds on Romer's model, and Jones's adaption. Following that, we translate the conceptual framework into our empirical approach. We then present the results. We follow that with a discussion of alternative explanations for declining research productivity. Finally, we discuss the implications for public policy.

1. Conceptual framework

Romer's (1990) theory of endogenous growth generates an equilibrium of balanced economic growth, in which the rate of knowledge growth drives output growth. Moreover, it yields a specific “scale effects prediction” that growth, g, should be proportional to the level of research labor, H_A (Eq. (1)). Recent U.S. experience, first documented by Jones (1995), is at odds with that prediction. In particular, while R&D spending has been rising, GDP growth has been declining. This decline in research productivity at the macro level has been replicated at the micro level across a number of domains (BJVW).

$$g = \frac{\dot{Y}}{Y} = \frac{\dot{A}}{A} = \delta\, H_{\mathrm{A}} \tag{1}$$

Jones (1995), who first documented the prediction's failure, proposed that the problem lies with the prediction itself. In particular, Romer's equilibrium of balanced growth is based in part on an assumption of constant returns to scale for each input in the knowledge production function, which generates new ideas (knowledge) from the stock of prior ideas, A_t, and research labor, H_A (Eq. (2)).

$$\dot{A} = \delta H_A A_t \tag{2}$$

Jones allowed for the possibility of non-constant returns by introducing λ, the elasticity of research labor and φ, the elasticity of the knowledge stock (Eq. (3)). Jones argued that there may be duplication in the R&D process, “externalities,” such that λ is <1. He further argued that the discovery of new ideas may be decreasing in the level of knowledge, “fishing out,” such that φ is <1. Only when λ and φ are both equal to 1 does Eq. (3) reduce to Romer's knowledge production function.

$$\dot{A} = \delta H_A^{\lambda} A^{\varnothing} \tag{3}$$

If the knowledge production function is captured by Eq. (3), then the corresponding growth equation is given by Eq. (4). Accordingly, if Jones is correct that λ and φ are both <1, then growth from R&D converges to zero. Conversely, if φ is greater than or equal to 1, there is no fishing out, and growth from R&D continues in perpetuity.

$$g = \delta H_A^{\lambda} A^{\varnothing - 1} \tag{4}$$

We empirically tested Jones's production function (Eq. (3)) to determine if it exhibited fishing out for U.S. firms. This is an economically meaningful test, because 71 % of R&D in the United States is conducted by firms.

2. Empirical specifications and data

2.1. Test of the knowledge production function

We utilized the *research productivity* measure in BJVW to test Jones's knowledge production function. BJVW constructed *research productivity* α by dividing growth by the number of researchers used to achieve that growth (Eq. (5)).[1]

[1] Equations 5 and 14 in Bloom et al. (2020).

$$\alpha = \left(\frac{\dot{A}_t}{A_t}\right) \Big/ H_A \tag{5}$$

To test Eq. (4), we needed to distinguish between *research productivity*, α, in BJVW and the productivity parameter, δ in Eqs. (1)–(4). To do so, we rearranged the terms in Eq. (5) to express growth in terms of *research productivity*: $g = \alpha H_A$, and then substituted for g in Eq. (4), and divided both sides by H_A (Eq. (6)):

$$\alpha = \delta H_A^{\lambda-1} A^{\varnothing-1} \tag{6}$$

We tested Eq. (6) using the following empirical specification (Eq. (7)):

$$ln(\text{research productivity}_{\text{it}}) = \beta_0 + \beta_1 ln(\text{KnowledgeStock}_{\text{it}}) + \beta_2 ln(\text{Scientists}_{\text{Ait}}) + \gamma_{\text{i}} + \delta_{\text{t}} + \varepsilon_{\text{it}} \tag{7}$$

which includes industry fixed effects, γ_i, and year effects, δ_t. We used industry fixed effects rather than firm fixed effects, because firm fixed effects subsume selection effects and the treatment effects of research productivity.

Once we obtained coefficient estimates of Eq. (7) for the *KnowledgeStock*, A_{it}, and *Scientists*, H_{Ait}, we converted them to values in the production function as follows: $\varphi = \beta_1 + 1$ and $\lambda = \beta_2 + 1$. If the knowledge stock is subject to fishing out, we would expect $-1 < \beta_1 < 0$.

2.2. Testing across the research productivity distribution

The observed declines in firms' *research productivity* in BJVW pertains to the mean. If knowledge production is subject to fishing out, the decline should be evident across the entire distribution of firms' *research productivity*, not merely the mean. In other words, if ideas are truly being fished out, then the quality of the best idea each year should, on average, be worse than the best idea in the prior year. Accordingly, we examined the time trend in maximum research productivity as a supplemental test. (We ignored the minimum because it is bounded by 0.) As in the main empirical model, we included industry fixed effects, except when looking across the entire economy, where there was only one observation per year. If ideas are getting harder to find, we expected the coefficient of β_1 in Eq. (8) to be negative and significant.

$$\text{Maximum Research Productivity}_{\text{it}} = \beta_0 + \beta_1 \text{ year}_t + \gamma i + \varepsilon \text{ it} \tag{8}$$

2.3. Data

Data on firms' *effective research* (which we labeled *scientists*) and *research productivity* came from data and codes made available by BJVW through the *American Economic Review* data-availability policy.[2] Their source data is the Compustat Fundamentals Dataset, which provides financial data on firms publicly traded on U.S. exchanges. BJVW formed four separate measures of decadal-average *research productivity*, which differed in the form of growth: sales revenue, market capitalization, employment, and revenue labor productivity (sales/employment). The use of firm-decade rather than firm-year was intended to smooth out fluctuations.

BJVW formed *effective research* in each firm-year by dividing firms' R&D expenditures by the mean U.S. wages of male workers with four or more years of college. This restricted the sample to firms reporting non-zero R&D (approximately 40 % of publicly traded firms). In addition, BJVW discarded (a) firms with less than three annual growth observations in a given decade, (b) firm-decades in which growth was negative, since negative growth is not a consequence of R&D, and (c) firms lacking data on both growth and R&D for two consecutive decades.

While BJVW's use of the wages of male college graduates seems to be an odd (and possibly controversial) choice, it stems from the fact that science and engineering (S&E) wage data are only available for three years (1995, 2003, 2017), whereas college graduate (BS) wage data are available annually. Comparison of S&E wages to male BS wage data for the three years in which both are available indicates that male BS wages closely approximate mean S&E wages. Male S&E wages are well above male BS wages, while female S&E wages are well below male BS wages. However, in robustness checks, we employed two alternative measures of *effective research*: (a) *advanced degrees*: firm R&D divided by the U.S mean wage for employees with advanced degrees, and (b) *inventors:* the number of inventors on all patents filed by a firm in a given year.

Effective research enters Eq. (7) both directly as ln*(scientists)* and indirectly as the denominator for firm *research productivity* (growth/*scientists*). As mentioned previously, BJVW used four alternative forms of firm growth. Since their results are comparable across the four growth measures, we utilized only one form of growth, sales revenue. Our only deviation from the BJVW data is that they utilized decadal summaries of *research productivity*. This was sufficient for BJVW because they merely reported decadal declines. However, our econometrics require greater resolution, so we interrupted their code to extract firm-year measures of *revenue growth* and *effective research.* We then created annual *research productivity* by dividing average *revenue growth* over the prior ten years by median *effective research* over the same period.

The biggest data challenge in testing the knowledge production function was constructing measures for the knowledge stock, A. The convention is to accumulate and depreciate either R&D expenditures or patent counts. We cannot accumulate R&D because that would confound the knowledge stock with *effective research*. Accordingly, we utilized firms' patents. We constructed three proxies for firms' knowledge stocks using the U.S. patent dataset released by Kogan et al. (2017) and updated by Stoffman et al. (2020). To form each of the proxies, we began by counting all patent applications in each cooperative patent classifications group per year. We then constructed the yearly stock of patents per technology group using a depreciation rate of 20 %. This is in line with prior studies, indicating that estimates are insensitive to depreciation rates from 8 % to 25 % (Hall et al., 2010).[3] Next, for each patent, p, we formed its knowledge stock, $KSPat_p$, as the average of all groups' knowledge stocks in the patent, weighted by the number of times each group was used to classify the patent.[4]

We then constructed two proxies from $KSPat_p$ as follows.*PatentKnowledgeStoc* is the average knowledge stock per patent filed by firm i in year t: $(\Sigma_{pit} KSPat_p)/\#\text{patents}_{it}$. *CitationKnowledgeStock*$_{it}$ averages out the $KSPat_p$ of its citations. *PatentKnowledgeStock*$_{it}$ captures the knowledge a firm is creating, while *CitationKnowledgeStock*$_{it}$ captures the knowledge upon which a firm is building. Finally, we constructed a third measure, *CitationValue*$_{it,}$ which captures the value of all patents cited by the focal firm in a year, as estimated by Kogan et al. (2017). All measures of knowledge stock use a rolling ten-year window to be consistent with the *research productivity* measure.

To merge the BJVW data with the knowledge stock measures, we followed Kogan et al. (2017) in using the CRSP-COMPUSTAT link table in the CRSP/Compustat Merged Database. The final dataset comprises all U.S.-traded firms that conducted R&D and filed patents from 1983 to 2015, with sufficient observations to form *research productivity* (a minimum of three years of data, to be consistent with BJVW). These data are summarized in Table 1. The table indicates that, on average, firms have \$5.3 billion in revenues and invest \$214 million in R&D each year, which translates to 1833 effective researchers. Note that the restrictions mentioned above reduce the set of firm-year observations for firms

[2] https://www.aeaweb.org/articles?id=10.1257/aer.20180338.

[3] The results are robust to depreciation rates of 40 %, 60 %, and 80 %, likely because stocks and flows are econometrically equivalent in steady-state (Griliches and Mairesse, 1984)

[4] The patent is classified at the subgroup level. Thus, a patent can be classified into multiple subgroups that share the same group.

Table 1
Descriptive statistics.

Variable	Obs.	Mean	St. Dev.	Min	Max
Ln(Research Productivity)	37,971	−7.53	2.50	−18.82	0.93
Research Productivity	37,971	0.01	0.03	0.00	2.55
Ln(Research Productivity_inventors)	36,927	−4.40	2.14	−15.96	0.94
Research Productivity_inventors	36,927	0.05	0.10	0.00	2.56
Ln(Scientists)	37,971	5.25	2.00	−1.90	11.74
Scientists	37,971	1833.40	7630.04	0.15	1.3e+05
Ln(Inventors)	36,927	2.11	1.53	0.00	8.95
Inventors	36,927	51.47	241.75	1.00	7685.10
Ln(Patent Knowledge Stock)	35,543	8.31	1.00	3.09	11.38
Patent Knowledge Stock	35,762	12,069.39	14,465.03	19.91	1.7e+05
Ln(Citation Knowledge Stock)	35,762	8.80	1.14	2.99	12.04
Citation Knowledge Stock	35,543	6529.53	7109.57	22.00	87,322.38
Ln(Citations Total Value)	36,318	5.32	2.46	−5.61	14.76
Citations Total Value	36,927	5204.70	39,128.53	0.00	2.6e+06
Revenue	22,269	5264.49	23,478.47	0.01	4.8e+05
Ln(Revenue)	22,269	5.67	2.49	−4.96	13.07
R&D	22,267	214.34	817.82	0.00	12,540.00
Ln(R&D)	22,267	3.15	2.07	−6.91	9.44

Notes: Research productivity is calculated as the average revenue growth over the prior ten years by median effective research over the same period. Research productivity_inventors is calculated as the average revenue growth over the prior ten years by average number of inventors over the same period. Scientists is the average R&D divided by U.S. mean wage for employees with advanced degrees over the prior ten years. Inventors is equal to the average number of distinct inventors over the prior ten years. Patent knowledge stock is the average knowledge stock per firm in the current year. Citation knowledge stock is the average knowledge stock contained in the firm's citations in the current year. Citations total value is the value of all patents cited by the focal firm in the current year. Revenue and R&D are in $millions.

conducting R&D by roughly two thirds—from 103,000 firm-year observations to 38,000. The firms in our dataset represent a broad swath of the economy, as shown in Fig. 2. They are larger on average than the broader set of firms. They average $5.3 billion in annual revenues versus $2.3 billion. However, the R&D intensity (R&D/revenues) is similar across the two groups. The $214 million average R&D represents a 4.0 % R&D intensity compared to the $89 million average R&D and 3.9 % R&D intensity on the broader dataset.

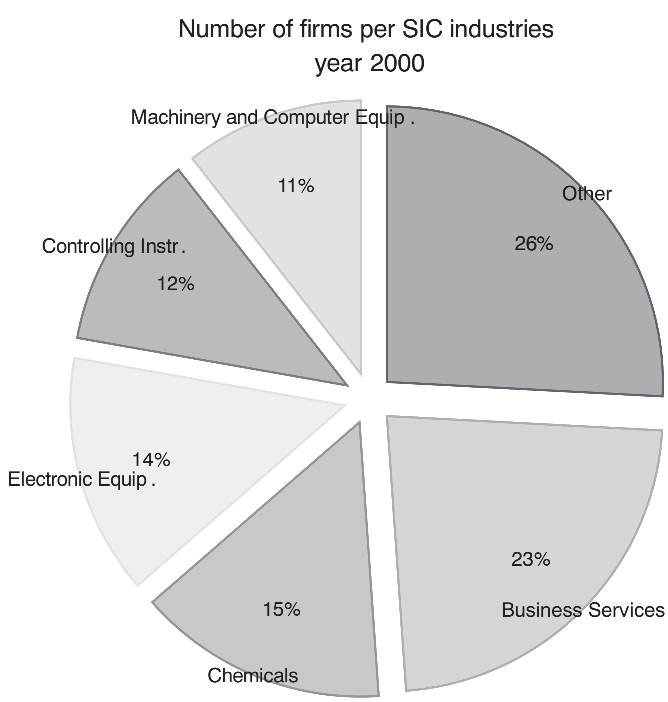

Fig. 2. Industry distribution of firm sample
Notes: In this figure we characterize the distribution of the sample across industry for 2020. Note that this distribution differs from the universe of U.S. firms in three ways: (a) all firms in the sample are publicly traded (approximately one third of U.S. employment), (b) all firms in the sample conduct R&D (approximately 40 % of publicly traded firms), (c) all firms in the sample file patents (approximately 50 % of firms who conduct R&D).

3. Results

3.1. Test of the knowledge production function

Table 2 presents the test results of the knowledge production function (Eq. (7)). Looking first at fishing out, the coefficient estimates for the knowledge stock measures, β_1, are always positive and significant. The estimates for β_1 (0.02 to 0.24) imply values for φ of 1.02 to 1.24 ($\varphi = \beta_1 + 1$). Thus, we failed to find support for fishing out. Instead, we found mild evidence of increasing returns to the knowledge stock. A 10 % increase in the knowledge stock increases new knowledge creation by 10.2 to 12.4 %.

Looking next at externalities, the coefficient β_2 on ln*(scientists)* is negative and significant in all models. The estimate for β_2 (−1.1) implies

Table 2
Test of the knowledge production function

	Ln(Research Productivity)		
	(1)	(2)	(3)
Log(Scientists)	−1.10***	−1.10***	−1.11***
	(0.01)	(0.01)	(0.01)
Log(CitationKS)	0.18***		
	(0.02)		
Log(PatentKS)		0.24***	
		(0.03)	
Log(Citation value)			0.02***
			(0.01)
Constant	−3.41***	−3.74***	−1.84***
	(0.19)	(0.22)	(0.06)
R-squared	0.80	0.80	0.80
Observations	35,761	35,542	36,318
Industry FE	SIC 4	SIC 4	SIC 4
Year FE	Yes	Yes	Yes

Notes: In this table we examine the transformed knowledge production function to estimate the elasticity of the knowledge stock, for three separate measures of the knowledge stock. Column 1 uses CitationKnowledgeStock, Column 2 uses PatentKnowledgeStock, and Column 3 uses citation value. We include industry (defined by 4-digit SIC code) fixed effects, and year effects. Standard errors clustered at the firm level are in parentheses.
*** $p < 0.01$.
** $p < 0.05$.
* $p < 0.1$.

a value for λ of −0.1 ($\lambda = \beta_2 + 1$). This suggests that a 10 % increase in research labor would reduce new knowledge creation by 1 %, which seems implausible. It is likely an artifact of BJVW's use of *scientists* to form *research productivity*.

Robustness checks with two alternative measures of effective research (advanced degrees and inventors) match those in Table 2. The coefficient estimates for β_1 range from 0.03 to 0.20, and the estimates for β_2 range from −1.08 to −1.20.

3.2. Testing decline across the research productivity distribution

The results for test of the trend in *maximum research productivity* are presented in Table 3. For reference, Model 1 presents the time trend for mean firm *research productivity*. The coefficient on year in Model 1 is negative, reflecting the decline in mean research productivity observed by Jones (1995). Looking next at *maximum research productivity* across the entire economy (Model 2), the coefficient estimate on year (0.146) is positive and significant. This means that the most productive firm in each year grows revenues 14.6 % faster than the most productive firm in the prior year. Thus, results with this supplementary test are consistent with those from the formal test in Table 2. We found no evidence of fishing out. Rather, the results suggested that there are positive spillovers in knowledge production.

Because this result was unanticipated, we explored it further to see if by narrowing domains, we could mimic the results in BJVW, who looked within domains. Models 3 through 6 examine increasingly narrowed domains, where domains are defined as the firm's primary SIC code. Model 3 estimates the trend in *maximum research productivity* within 1-digit industries, Model 4 within 2-digit industries, Model 5 within 3-digit industries, and Model 6 within 4-digit industries. Across Models 3 to 6, coefficient estimates are positive and significant. However, their magnitude decreases as domains are narrowed. For example, while *maximum research productivity* grows 15 % per year across the economy, it only grows 2 % per year within 1-digit industries, 0.4 % within 2-digit industries, 0.2 % within 3-digit industries and 0.1 % within 4-digit industries. This implies that opportunity is greater across industries than within them.

While the results in Table 3 utilize SIC definitions of industry, we obtained similar results when using NAICS definitions of industry. For either industry definition, *maximum research productivity* is always increasing over time, but it increases at lower rates as industry definition is narrowed. Note that BJVW's measure of *research productivity* represented a firm mean—growth divided by *effective research*. When we replicated the test using a measure of marginal productivity (the firm's output elasticity of R&D), the coefficient on year became negative for 2-digit industries and became increasingly negative as industry definition was narrowed. Thus, when using a marginal measure, *maximum research productivity* is decreasing within industries over time.

Taken together, the results for *maximum research productivity* suggest that opportunity within industries decays over time. However, as it does, it appears that firms create new industries with greater opportunity. This matches Schumpeter's (1942) notion of creative destruction. A classic example is automobiles replacing horses and buggies. A more recent example is personal computers replacing electronic typewriters. At the U.S. peak, the installed base of electronic typewriters reached 10 million units. By contrast, the installed base of the PCs replacing them reached 316 million units in 2008. Thus, not only did the PC replace the electronic typewriter, but it also seems to have created much greater value.

While this finding of decreasing opportunity within industry, yet increasing opportunity across industries is interesting in and of itself, the goal of this exercise was a robustness check of the main results in Table 2. As in the main test, here again, we failed to find support for fishing out.

3.3. Summary

In summary, for both a structural test of the knowledge production function, and an alternative test across the research productivity distribution, we failed to find evidence of fishing out. In fact, the elasticity of the knowledge stock in generating new knowledge was positive and significant for all measures, suggesting there are positive spillovers in the production of knowledge, rather than fishing out. This result is reinforced by the observation that *maximum research productivity* is increasing over time, which we would expect if there were positive spillovers in the production of knowledge.

4. Alternative explanations for declining research productivity

While we failed to find evidence of fishing out, there was substantial evidence that aggregate *research productivity* has declined. What, other than fishing out, might explain that? We propose that the research productivity parameter, δ, in Eq. (1), is itself declining. Indeed, the year effects in estimating Eq. (7) continue to exhibit decline after modeling the contributions from knowledge and research labor (Fig. 3). We further propose that the decline in δ may be due to contingent factors at both the macro and micro levels. We discuss one factor at each level. At the economy level, there is growing imbalance between research and development. At the firm level, there is evidence of deterioration in R&D practices. Note these two factors are not intended to be exhaustive. Rather their role, together with Fig. 2, is to provide suggestive evidence

Table 3
Assessing the trend in maximum research productivity.

	Mean a	Maximum Annual Research Productivity, aMAX				
	(1)	(2)	(3)	(4)	(5)	(6)
Year	−0.000**	0.146***	0.020**	0.004**	0.002***	0.001***
	(0.000)	(0.053)	(0.008)	(0.002)	(0.001)	(0.000)
Constant	0.26**	−288.83**	−38.83**	−8.37**	−3.66***	−2.40***
	(0.117)	(105.957)	(15.400)	(3.354)	(1.219)	(0.802)
Observations	55,988	33	297	1633	5238	8082
R-squared	0	0.16	0.184	0.186	0.144	0.141
Industry FE	NO	NO	sic1	sic2	sic3	sic4

Notes: In this table, we examine the trend in *maximum research productivity*, defined as the maximum observed value of *research productivity* in a given year. For reference, Column 1 presents the trend in mean *research productivity.* Column 2 estimates the trend in *maximum research productivity* across the entire economy in each year, Column 3 captures the mean trend within 1-digit industries, Column 4 captures the mean trend within 2-digit industries, Column 5 captures the mean trend within 3-digit industries, and Column 6 captures the mean trend within 4-digit industries. Standard errors are clustered at the industry level. Robust standard errors are in parentheses.

*** $p < 0.01$.

** $p < 0.05$.

* $p < 0.1$.

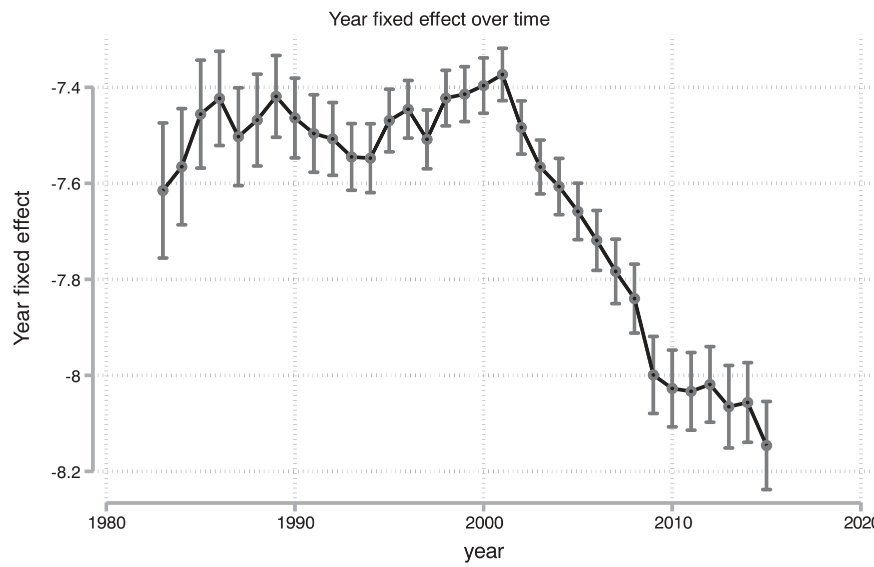

Fig. 3. Year effects from estimating research productivity
Notes: In this figure we plot the year effects obtained from estimating Eq. (7).

that contingent factors (rather than fishing out) are responsible for the declines in research productivity and economic growth.

4.1. Increasing imbalance between research and development

For research (R) to generate growth, it must be commercialized or otherwise diffused. This requires development (D). Development expenditures within a firm typically dwarf the associated research expenditures. For example, in the case of pharmaceuticals, total expenditures for clinical development are approximately twice those for pre-clinical research. For any given drug, the ratio is even more pronounced—development expenditures are 16 times pre-clinical expenditures (Mestre-Ferrandiz et al., 2012). The difference between total expenditures and drug-specific expenditures arises from the need to carry a portfolio of pre-clinical drugs to compensate for ones that do not survive to clinical stages. Historically U.S. federal R&D expenditures have mimicked the ratio of R to D within firms. In 1961, federal research expenditures comprised 0.51 % of U.S. GDP, while development expenditures comprised 1.23 % of GDP. Since that time, however, the relative proportions have reversed. Research funding has remained near its prior level (0.38 % of GDP in 2018), while development funding has fallen to roughly half that (0.22 % of GDP) (Fig. 4). Moreover, this decline in federal development funding has coincided with the decline in aggregate research productivity (Fig. 5). Thus, the shift in federal allocation of R&D investment may be contributing to the decline in U.S. research productivity. What seems likely, though we cannot observe it, is that the federal shift in R&D allocation, may have generated “excess research”—meaning that the number of inventions generated by universities and federal labs may now exceed industrial capacity to develop them.

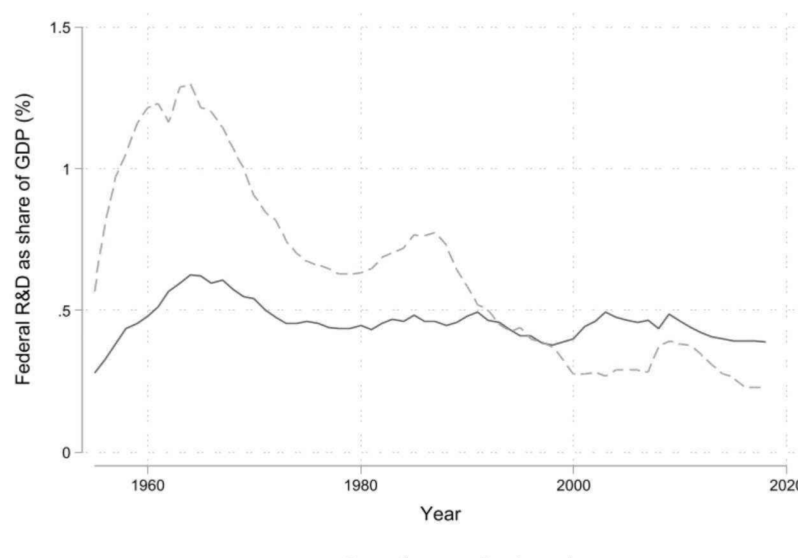

Fig. 4. Evolution in allocation of federally funded R&D
Notes: In this figure we plot allocation of federally funded R&D to all performers on research (basic plus applied), and on development from 1960 to 2018. *Source: National Patterns of R&D Resources: 2017–18 | Detailed Statistical Tables| NSF 20–307 | January 08, 2020.*

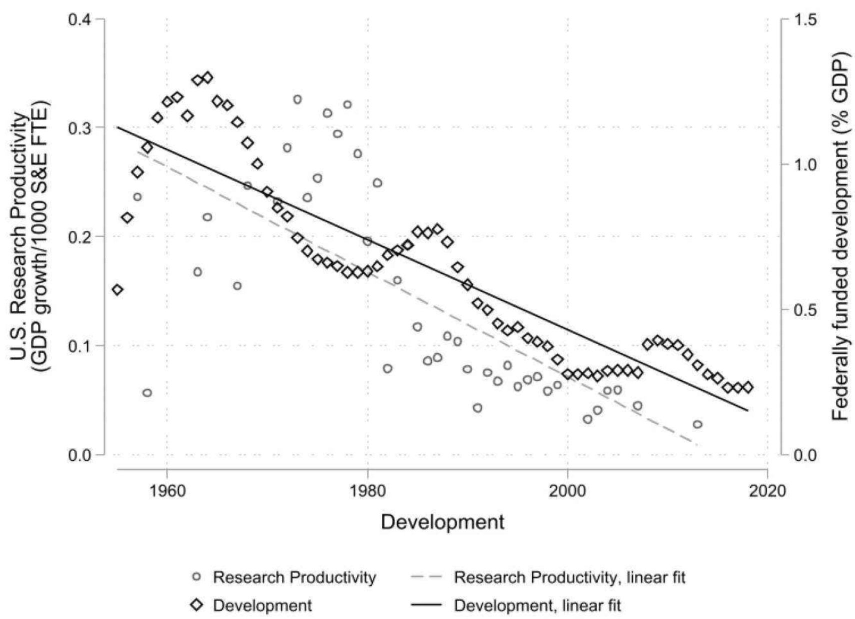

Fig. 5. Relationship between federally funded development and research productivity
Notes: In this figure we plot federally funded development to all performers from 1950 to 2018. We compare this to aggregate research productivity measured as GDP growth (constant dollars/constant dollars) divided by research labor (measured as 1000 full-time equivalent scientists and engineers). *Source: National Patterns of R&D Resources: 2017–18 | Detailed Statistical Tables| NSF 20–307 | January 08, 2020.*

4.2. Deterioration in research practices

Just as management practices affect firms' total factor productivity (TFP) (Bloom and van Reenen, 2010; Knott, 1996), R&D practices affect firms' research productivity. However, while firms tend to evolve toward better management practices, there is emerging evidence that they evolve toward inferior R&D practices. The divergence may stem from the fact that R&D outcomes have long lags, so it is harder to understand the relationship between practices and performance for R&D than it is for other functions.

Three documented trends toward inferior R&D practices are (a) a 34 % decrease in the level of R&D centralization (Argyres and Silverman, 2004; Arora et al., 2011; Cummings, 2018), (b) a six-fold increase in the intensity of R&D outsourcing (Knott, 2020), and (c) a 67 % increase in the rate at which CEOs are hired from outside the firm (Cummings and Knott, 2018). Each of these practice changes is associated with lower research productivity.

With regard to decentralized R&D, Argyres and Silverman (2004), and Arora et al. (2011), both found that it leads to patents that are less broad and less impactful. Cummings (2018) further found that decentralization is associated with significantly lower patent intensity and R&D productivity.

With regard to outsourced R&D—R&D funded by one company, but performed by another—Knott (2020) found the mean output elasticity of outsourced R&D was precisely estimated at zero, versus a firm mean of 0.13 for internal R&D. This means that a 10 % increase in internal R&D increases firm revenues by 1.3 % on average, whereas outsourced R&D has no impact on firm revenues.

With regard to outside CEOs, Cummings and Knott (2018) found that R&D productivity decays with each year of an outside CEO's tenure. They further found that there is less decay if the CEO is from the same industry, suggesting that the decay stems from lack of expertise. Kluppel and Cummings (2019) extended this work to examine whether outsiders without the requisite expertise change the direction of R&D in detrimental ways. They found instead that outside CEOs seem to do the opposite. They maintain the existing technological trajectory, which itself is detrimental if technology is moving in a different direction.

Taken together, evidence at both the macro and micro levels suggests that the organization and conduct of R&D has changed over the period of declining research productivity. Moreover, there is theory and evidence to suggest that these changes are associated with lower R&D productivity.

5. Discussion

Innovation is the primary source of economic growth, yet it is failing to deliver that growth in the United States, because research productivity has been declining. Scientific labor has been increasing, while GDP growth has been decreasing. The leading explanation for the decline is that there are decreasing returns to the knowledge stock, or “fishing out” (Jones, 1995). If correct, then the United States will need to spend increasing amounts of R&D to maintain even the current level of growth. Ultimately R&D will produce no growth. In other words, there would be no point in either public or private investment in R&D.

We tested the fishing out explanation by empirically characterizing the knowledge production function for U.S. firms. We failed to find support for fishing out. Instead, we found that there were increasing returns to the knowledge stock. In addition to this formal test, we conducted a supplemental test by examining trends in *maximum research productivity*. Here too, we found that *maximum research productivity* has been increasing over time, which we would expect if there are increasing returns to the knowledge stock.

While our tests failed to find evidence of fishing out, they do support the observation that firms' research productivity has declined. Given that fishing out cannot explain the decline, we proposed that the research productivity parameter, δ, in Romer's knowledge production function, is itself declining. Indeed, research productivity year effects continued to exhibit decline after we modeled contributions from the knowledge stock and research labor. We identified contingent factors that may contribute to the decline in δ, and we provided preliminary evidence of their role. At the economy level, we documented an increasing imbalance between research and development, and we showed that the decline in federally funded development coincides with the decline in aggregate research productivity. At the firm level, we documented three trends in R&D practices: increased decentralization of R&D, increased outsourcing of R&D, and increased use of outside CEOs. We then reviewed the empirical literature demonstrating the correlation of these newer practices with lower R&D productivity. Thus, it appears plausible that contingent factors are responsible for at least some of the decline in research productivity.

In summary, our results failed to find support for fishing out. Rather we found positive spillovers of the knowledge stock in the production of new knowledge. Thus, we do not need to be concerned that R&D is getting harder, or accordingly that growth from R&D will decline to zero. Rather, these results support Romer's original form of the knowledge production function, as well as its expectation that R&D investment can generate growth in perpetuity.

However, and this is critically important, both the government and firms need to recognize that U.S. research productivity has declined dramatically. It had declined 70 % at the time of Jones's first documentation (Jones, 1995). It has declined again by the same amount in the period since then, as Fig. 5 shows. Therefore, to avoid expending increasing amounts of R&D to maintain even the current levels of growth, we need to identify and ameliorate contingent factors contributing to the decline. We provide evidence of such factors at the macro and micro levels.

While our analysis pertains to firms, the findings have important implications for public policy. First, the vast majority of R&D is performed by industry. This likely warrants greater attention to firms in federal innovation policy. The main existing policy instrument for firms is the R&D tax credit. Because of its structure, the tax credit rewards firms for increasing research, but not for increasing development. Thus, the tax credit is likely contributing to the problem of excess research associated with the federal shift from D to R.

Even ignoring an issue of excess research, aggregate data suggest that R&D spending has increased on average 0.014 % per year since 1978, when expressed as share of GDP. In absolute dollars, the increase is even more pronounced. Thus, it does not appear that firms need incentives to increase R&D investment. Rather, they need incentives to improve R&D productivity. One policy approach to accomplish this is tying the R&D tax credit to improvements in R&D productivity rather than increases in research spending.

The second implication of our firm-level analysis for public policy is that development, which is required to commercialize or diffuse inventions, is almost exclusively performed by industry. Therefore, if industry R&D is unproductive, then research done by universities and labs becomes de facto unproductive. It may continue to appear productive when measured by patents and publications, but if there is no capacity to develop those inventions, they cannot contribute to economic growth. To remedy this, policymakers could move toward restoring prior ratios of federal funding for development relative to research.

CRediT authorship contribution statement

Anne Marie Knott: Conceptualization, Methodology, Writing—Original Draft, Visualization

Leonardo Klüppel: Methodology, Software, Formal analysis, Data Curation, Writing—Reviewing and Editing, Visualization

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Argyres, N., Silverman, B., 2004. R&D, organization structure, and the development of corporate technological knowledge. Strateg. Manag. J. 25 (8–9), 929–958.

Arora, A., Belenzon, S., Rios, L., 2011. The Organization of R&D in American Corporations: The Determinants and Consequences of Decentralization. NBER Working Paper No. 17013.

Bloom, N., van Reenen, J., 2010. Why do management practices differ across firms and countries? J. Econ. Perspect. 24 (1), 203–224.

Bloom, N., Jones, C., van Reenen, J., Webb, M., 2020. Are ideas getting harder to find? Am. Econ. Rev. 110 (4), 1104–1144.

Clancy, M., 2018. Inventing by combining pre-existing technologies: patent evidence on learning and fishing out. Res. Policy 47 (1), 252–265.

Cummings, T., 2018. Innovation strategy and R&D structure: Internal and external contributors to R&D structure. Acad. Manag. Proc. 2018 (1).

Cummings, T., Knott, A.M., 2018. Outside CEOs and innovation. Strateg. Manag. J. 39 (8), 2095–2119.

Furman, J., Porter, M., Stern, S., 2002. The determinants of national innovative capacity. Res. Policy 31 (6), 899–933.

Griliches, Z., Mairesse, J., 1984. Productivity and R&D at the firm level. In: Griliches, Z. (Ed.), R&D, Patents, and Productivity. University of Chicago Press, Chicago, pp. 339–374.

Hall, B., Mairesse, J., Mohnen, P., 2010. Measuring the returns to R&D. In: Hall, B.H., Rosenberg, N. (Eds.), Handbook of the Economics of Innovation. North-Holland, Amsterdam.

Jones, C., 1995. R&D-based models of economic growth. J. Polit. Econ. 103 (4), 759–784.

Kluppel, L., Cummings, T., 2019. A New Direction or Lack of Direction? The Story of Innovation With Outsider CEOs. Working Paper.

Knott, A.M., 1996. Do Managers Matter? UCLA. Unpublished PhD Dissertation.

Knott, A.M., 2020. The Record on R&D Outsourcing. U.S. Census Bureau Center for Economic Studies Paper No. CES-WP-16-19.

Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N., 2017. Technological innovation, resource allocation, and growth. Q. J. Econ. 132 (2), 665–712.

Mestre-Ferrandiz, J., Sussex, J., Towse, A., 2012. The R&D cost of a new medicine. Office of Health Economics, London. https://www.ohe.org/publications/rd-cost-new-medicine.

Romer, P., 1990. Endogenous technological change. J. Polit. Econ. 98 (5), S71–S102.

Schumpeter, J., 1942. Capitalism, socialism and democracy. Harper & Row, New York.

Solow, R., 1957. Technical change and the aggregate production function. Rev. Econ. Stat. 39 (3), 320.

Stoffman, N., Woeppel, M., Yavuz, M.D., 2020. Small innovators: no risk, no return. Working Paper. https://ssrn.com/abstract=3291471.