Staff Paper P22-06 April 2022

Staff Paper Series

Are Ideas *Really* Getting Harder to Find?

by

Julian M. Alston and Philip G. Pardey

Department of APPLIED ECONOMICS

INSTEPP INTERNATIONAL SCIENCE & TECHNOLOGY PRACTICE & POLICY

College of Food, Agricultural and Natural Resource Sciences

UNIVERSITY OF MINNESOTA

Are Ideas *Really* Getting Harder to Find?

April 18, 2022

Julian M. Alston and Philip G. Pardey

Affiliations and Acknowledgements

Julian Alston is a Distinguished Professor in the Department of Agricultural and Resource Economics and Director of the Robert Mondavi Institute Center for Wine Economics at the University of California, Davis, and a member of the Giannini Foundation of Agricultural Economics. Philip Pardey is a Professor in the Department of Applied Economics, Director of Global Research Initiatives for CFANS, Director of the International Science and Technology Practice and Policy (InSTePP) center, and Director of the GEMS Informatics Center, all at the University of Minnesota.

The authors are especially grateful for the excellent research assistance provided by Connie Chan-Kang, Devin Serfas, and Shanchao Wang. We also give much thanks to Yuan Chai for providing access to and support in analyzing the novel wheat varietal evidence cited in this paper. The work for this project was partially supported by the Minnesota Agricultural Experiment Station (MIN-14-161); the University of Minnesota's GEMS Informatics Center; the USDA National Research Initiative; the California Agricultural Experiment Station; and the Giannini Foundation of Agricultural Economics.

ABSTRACT

Bloom et al. (2020) attribute the post-WWII slowdown in growth of U.S. TFP and other productivity measures to a decline in research productivity. A weakness in their approach is that the authors measure research productivity as the annual growth rate of industrial or economywide productivity divided by the number of researchers, contemporaneously. They give no consideration to the stock-flow relationships whereby current research effort gives rise to increments to a stock of depreciable knowledge and hence an evolving path of enhanced productivity over an extended but possibly finite future period. Using examples from agriculture, for which we have comparatively rich data, we revisit established ideas and evidence on links between research spending and productivity. On both conceptual and empirical grounds, we question whether the evidence supports the claim that a decline in productivity of researchers is responsible for the slowdown in productivity growth that has been observed, the large increases in numbers of scientists and in spending per scientist notwithstanding.

Keywords: agricultural productivity, slowdown, R&D lags, innovation

JEL Codes: D24, E23, O31, O47, Q16

Are Ideas *Really* Getting Harder to Find?

In a recent article, Bloom et al. (2020) ask "Are Ideas Getting Harder to Find?" Using what they present as a conventional idea-based growth model applied to data for various industries, products and firms, the authors conclude that "research productivity is falling sharply everywhere we look … [and indeed,] … ideas are getting harder and harder to find" (p. 1138). This important finding rests on a specific conception of ideas and the idea-generation process: a knowledge production function in which (1) economic growth is proportional to the number of researchers, and (2) units for ideas are defined such that a constant flow of new ideas leads to a constant percentage rate of growth in total factor productivity (TFP). Here, we challenge the claim that research productivity is falling and ideas are getting harder to find. In doing so we draw mainly on detailed evidence for U.S. agriculture, which was foremost among the examples used by Bloom et al. (2020) to make their case. We do not dispute that the number of researchers (or any other measure of research effort) has risen considerably relative to the rate of productivity growth in agriculture and in most other sectors of the economy in most places around the world. We do dispute the claim that this is evidence of a decline in researcher productivity, and we also question whether anyone can confidently say ideas are really becoming harder to find using such models and measures.

I. Concepts

To illustrate their thinking, Bloom et al. (2020) present an endogenous growth variant of a production function model, in which total labor supply is fixed and research spending (R), defined to be a fixed share of output ($s_t := R_t/Y_t$), leads to increases in the effective quantity of labor and thus productivity (A_t). Assuming the idea production function:

$$(1) \qquad \dot{A}_t = \alpha R_t,$$

Bloom et al. (2020) show how the long-run growth rate ($\dot{A}_t / A_t$) on a balanced growth path can be decomposed into the product of "research productivity," $\tilde{\alpha}_t$ and the number of "effective scientists" or "research effort," $\tilde{S}_t$, which they take to be equal to research spending deflated by the scientific (or by proxy, high-skilled worker) wage rate.

$$(2) \qquad \frac{\dot{A}_t}{A_t} = \tilde{\alpha}_t \tilde{S}_t.$$

In this conception, if research productivity is constant, endogenous growth requires that a constant number of researchers be able to generate constant exponential growth; if productivity growth is not keeping up with the number of scientists, then ideas must be getting harder to find.[1]

As the authors state, this decomposition highlights a "…stylized view of economic growth that emerges from idea-based growth models" (Bloom et al. 2020, p. 1104).[2] They acknowledge that interpreting the left-hand side of equation (2) as representing the flow of new ideas "…is just a convenient definition, and in some senses a more accurate title for this paper would be 'Is Exponential Growth Getting Harder to Achieve?' " (Bloom et al. 2020, p. 1109). We would have much less quarrel with the paper if that were its title and its claims were

[1] Jones (1995, 2005) raises concerns about the "scale effects" prediction of this model—that an increase in the *level* of resources devoted to R&D should increase the *growth rate* of the economy so much so that a growing number of researchers causes the growth rate of the economy to grow exponentially. He observes that: "The assumption embedded in the R&D equation that the growth rate of the economy is proportional to the level of resources devoted to R&D is obviously false" (Jones 1995, p. 762), and proposes a corrective amendment to the model. The work by Bloom et al. (2020) might also be construed as a way of reconciling the standard model with the facts.

[2] They refer to "classic" studies by Romer (1990) and Aghion and Howitt (1992) and "many recent" studies that use this approach, of which they cite several. They say "We follow much of the literature, including Aghion and Howitt (1992), Grossman and Helpman (1991), and Kortum (1997), and define ideas to be in units so that a constant flow of new ideas leads to constant exponential growth in *A*" (Bloom et al. 2020, p. 1108).

confined to that question. But these caveats notwithstanding the authors apply the predictions from equation (2) directly to data and make specific claims about "ideas" as though the definition is not merely "convenient." We question that practice. More importantly, the basic prediction—represented by equation (2)—is dubious because it rests on significantly unrealistic assumptions about the knowledge production process, in equation (1), as we explain next.

Models Linking R&D and Productivity

Hundreds of studies have been undertaken modeling the links between investments in R&D (or the quantity of R&D) and productivity both in agriculture and in other sectors of the economy.[3] Almost all these studies use models that imply a very different dynamic relationship between research investments and the time path of productivity compared with equation (2), in which today's research spending has an instantaneous and permanent effect on productivity. Instead, productivity is typically assumed explicitly or implicitly to depend on a stock of knowledge capital that is increased by current increments to knowledge (perhaps "ideas") that are the result of investments in R&D over past years (or decades) but also is reduced by decrements as knowledge in use becomes obsolete or depreciates for other reasons.[4] Below we sketch the elements of the predominant models used for industrial R&D and their counterparts for agricultural R&D, and contrast them with that of Bloom et al. (2020).

The typical model imposes a constant elasticity form for the relationship between *TFP* and the knowledge stock in use (K):

[3] Rao et al. (2019) review 492 studies reporting 3,426 estimates of rates of return to agricultural R&D. Serfas et al. (2022) review 128 studies reporting 1,464 estimates of rates of return to other industrial R&D.

[4] Pardey et al. (2010) elaborate on the relevant arguments and supporting evidence for this conception of the (agricultural) research-innovation-adoption-disadoption process, as first laid out by Evenson (1967).

(3) $$TFP_t = \beta_0 K_t^{\beta_1}.$$

The knowledge stock itself is represented by a weighted sum of past investments in R&D:

(4) $$K_t = \sum_{n=0}^{\infty} w_n R_{t-n}.$$

The lag weights (w) can be also interpreted as measuring the contribution to the future knowledge stock from research investments in the current year:

(5) $$w_n = \frac{dK_{t+n}}{dR_t}.$$

They are typically held constant, implicitly assuming constant research productivity in this sense.

Considering both the R&D lag (it takes time to develop useful technologies based on the new "knowledge" resulting from research, which itself takes time) and the adoption process (as the new technology takes time to penetrate the market, and more so for technologies embodied in durable capital), we would expect the lag weights to start at zero but eventually to rise for a time with the length of the lag. The lag weights might stay at the eventual maximum forever, but in the vast majority of models at some point the lag weights begin to decline. This decline reflects depreciation of the stock of knowledge in use—either because the resulting innovations have become less effective in a changing physical and economic environment, or are supplanted because they have become less effective or have been made obsolete by subsequent innovations. As a practical example, co-evolving pests and diseases and changes in climate are specific causes of declining effectiveness of particular agricultural innovations (e.g., new pest-resistant varieties or pesticides), giving rise to a demand for maintenance agricultural research to prevent yields from falling (see, e.g., Olmstead and Rhode 2002).

Archetype R&D Lag Models

Two R&D lag distribution models, reflecting these ideas, have come to predominate in applications to agriculture: the 35-year trapezoidal model from Huffman and Evenson (1993) and the 50-year gamma distribution model from Alston et al. (2011). Both models were developed in studies of the effects of aggregate public investments in agricultural R&D on U.S. farm productivity, so the lag distributions reflect the aggregation of diverse innovations, some entailing very short lag processes and others with much longer lags, all measured across decades of data. These models have in common a finite overall R&D lag process comprising an initial gestation period before R&D has any impact on productivity; a period of rising impact; and a period of falling impact, eventually to zero. Pardey et al. (2010) provide detailed documentation justifying the use of models of this nature to represent the links between agricultural R&D and productivity. Notably, in these models many years elapse before R&D has its peak impact on productivity reflecting the fact that the processes of research, creation of knowledge, and the development and adoption of innovations all can take considerable time.

In applications to other industries, as described by Hall et al. (2010, p. 1,047) "…the workhorse of R&D stock estimation remains the perpetual inventory model…" in which research provides increments to a geometrically declining knowledge stock, where δ is the depreciation rate and g is a gestation lag:

$$K_t = K_{t-1}(1-\delta) + R_{t-g} = \sum_{n=0}^{\infty}(1-\delta)^n R_{t-n-g} \quad (6)$$

As typically used, this model allows little or no time for the process of knowledge creation and adoption: research has its maximum impact on productivity immediately (with zero gestation lag) or almost immediately (with a two-year gestation lag as suggested by Li and Hall 2020).[5] Thereafter the lag weights decline monotonically, geometrically, and typically rapidly given high assumed rates of knowledge depreciation.[6]

Critical Differences

The model employed by Bloom et al. (2020) differs from these mainstream empirical models in two ways. First, research spending has an immediate impact on productivity—the R&D process takes no time and maximum adoption is achieved instantaneously, as in most studies of industrial R&D but unlike any studies of agricultural R&D.[7] Second, the impact on productivity is permanent and—unlike almost any empirical studies of either industrial or agricultural R&D but in keeping with the prevalent practice in endogenous growth models—no allowance is made for obsolescence or depreciation of knowledge.

In contrast, the archetype empirical model of agricultural R&D entails a finite overall lag—eventually (in 50 years' time according to Alston et al. 2011, or sooner according to most others) any innovations resulting from today's research investment will cease to contribute to production and productivity.[8] Consequently, holding research productivity (i.e., the lag weights)

[5] This remains the predominant practice 30 years after Griliches (1992, pp. S41–42) declared authoritatively: "… the more or less contemporaneous timing of such effects is just not possible."

[6] Serfas et al. (2022) compiled 1,464 estimates of rates of return from 128 studies of industrial R&D. Of those 1,464 estimates, 97.3% were based on a perpetual inventory model; 88.2%, did not allow for any gestation lag, 64.4% used a knowledge depreciation rate of δ = 15% per year, and another 4.5% used a δ > 15% per year.

[7] Jones and Summers (2020) begin with a model in the same spirit as Bloom et al. (2020) and examine several reasons why the implied benefit-cost ratio may be too high, including a misspecified R&D lag model. They say "The above baseline assumes that the payoff from R&D investments occurs immediately. Yet there may be substantive delays in receiving the fruits of R&D investments" (Jones and Summers p. 13). "Aggregating across the different types of research, a middle-of-the-road delay estimate may be 6.5 years…" (Jones and Summers p. 14).

[8] In the archetype industrial R&D models, an infinitely long geometric lag is combined with a high depreciation rate (typically 15% per year) such that the effective overall lag length is only a decade or so before past research ceases to materially affect current productivity.

constant, if research spending is held constant, eventually a steady state will be reached in which all research is effectively maintenance research. In this steady state, the additions to the knowledge stock resulting from new innovations will be just sufficient to offset depreciation, the usable knowledge stock will be constant and so too will be productivity. Likewise, in the perpetual inventory model with geometrically declining lag weights implied by a depreciation rate of δ, in long-run equilibrium a constant flow of research investments implies a constant knowledge stock: $R_t = \delta K_t$.[9]

The key implication is potentially surprising: a constant flow of research investments implies *zero* productivity growth. But this result follows clearly from the empirical models that have widespread support in the literature; that it applies generally for models with finite lag structures can be seen by examination of equation (3). In contrast, in equation (1) a constant flow of research spending results in a constant *linear* rate of productivity growth, whereas imposing a constant research intensity (i.e., whereby a constant share of economic output is spent on research) implies a constant *proportional* rate of productivity growth, as in equation (2).

II. Evidence

Their distinctive modeling assumptions enable Bloom et al. (2020) to draw inferences about research productivity that cannot be drawn otherwise. We present a range of evidence against both those assumptions and their implications, as they pertain to agricultural R&D and productivity, but before doing that we note some concerns about their aggregative evidence.

[9] If the stream of research investments is growing at a constant rate, γ, then the knowledge stock will grow at the same rate: $K_t = \frac{R_t}{\delta - \gamma} = (1+\gamma)K_{t-1}$.

Aggregate Evidence

To begin to make their empirical case Bloom et al. (2020, p. 1111) plot a measure of the effective number of researchers (research "spending" divided by a measure of the nominal wage for high-skilled workers) for decades from the 1930s to the 2000s compared with measures of U.S. TFP growth (combining measures from BLS 2017 and Gordon 2016 in their Figure 1) and research productivity (the same proportional TFP growth expressed per effective researcher in their Figure 2). On this measure the effective number of researchers increases by a factor of 23 between the decade of the 1930s and the decade of the 2000s. We could not replicate that result exactly, but using a similar procedure with annual rather than decadal data we derived a corresponding measure that grew by a factor of 27; however, our measure grew at a diminishing rate (our plot is concave to the origin, rather than convex like that of Bloom et al. 2020).

In any event, on our reading these measures based on the stock of spending are conceptually flawed if the purpose is to measure the flow of effort. Their measure of the effective number of researchers is computed using a measure of a physical (or research) capital stock—"gross domestic investment in intellectual property products" from the National Income and Product Accounts, BEA (2017)—rather than a flow of current expenditures on research, divided by a wage rate for high-skilled workers. This ratio cannot be taken as a meaningful measure of the effective number of researchers.[10]

---

[10] Such an assumption would be questionable, even if they had a measure of the flow of research expenditure rather than the stock. Bloom et al. (2020, pp. 1112–1113) argue that "When the only input into ideas is researchers, deflating R&D expenditures by an average wage will recover a quality-adjusted quantity of researchers. In practice, R&D expenditures also include spending on capital goods and materials. As explained next, deflating by the nominal wage to get an 'effective number of researchers' that this research spending could hypothetically purchase remains a good way to proceed." But in empirical analysis of any economic activity it is perilous to assume labor is the only input or that labor is proportional to total expenditure divided by the wage rate, especially over very long periods of time or across different economic sectors.

We constructed measures of the quantity of annual research effort using the BEA gross domestic product series of total annual R&D spending for the period 1929–2018 (backcast to 1890) deflated by the InSTePP R&D deflator. The results are displayed in Figure 1a, along with our annual series of the rate of U.S. MFP growth (Pardey and Alston 2021, Table 2, p. 124). After the disruption from the Great Depression of the 1930s the rate of MFP growth fluctuates around a generally downwards trend. Similar patterns can be seen in the plot using less-detailed data (by the decades rather than annually) provided by Bloom et al. (2020, Figure 1, p. 1111). Meanwhile, our measures of the quantity of research effort increase by a factor of about 17 from the mid-1930s (1935) to the mid-2010s (2105). Substantive conceptual and quantitative differences in the measures notwithstanding, the picture is similar. That is, the annual rate of scientific effort has grown by an order of magnitude—10- or 20-fold—over a time period when the rate of MFP growth has trended down. No doubt, the rate of productivity growth per unit of scientific effort has gone down.

[Figure 1. *Aggregate on Growth, Research Effort, and Research Productivity*]

In their Figure 2, Bloom et al. (2020, p. 1111) plot the ratio of measures in their Figure 1, as a measure of "research productivity," along with their measure of research effort, now on a log scale. We created a counterpart figure as shown in our Figure 1b. On our measure, "research productivity" (moving average MFP growth per unit of real R&D spending) fell by a factor of 38 between 1935 and 2005, almost identical to the decline by a factor of 42 between the 1930s and the 2000s reported by Bloom et al. (2020). No doubt, the rate of productivity growth per unit of scientific effort has declined. But this purported measure of research productivity does not account properly for the temporal structure linking R&D and productivity.

Agricultural Knowledge Stocks and Productivity Patterns

Wang et al. (2022) use data on U.S. public agricultural R&D to illustrate the consequences of imposing various R&D lag distribution models for the estimates of knowledge stocks, models of multifactor productivity (MFP) growth as a function of investments in R&D, and the resulting estimates of elasticities and rates of return.[11] Among the models they estimated, as they explain, a gamma lag distribution model is preferred but, as estimated, it happens to be almost identical to the Huffman and Evenson (1993) trapezoidal lag distribution model. The geometric lag distribution model as proposed by Li and Hall (2018) yields superficially plausible estimates, even though it is wholly implausible in this application to data for U.S. agriculture—a context where numerous studies have documented the lengthy intrinsic lags between R&D spending and the generation of tangible, useful new ideas and the development and adoption of resulting innovations. In contrast, for these data the model of Bloom et al. (2020) does not even yield superficially plausible estimates—the estimated coefficient on the knowledge stock is statistically significant but has the wrong (i.e., negative) sign—and it performs relatively poorly as a statistical model. It is strongly rejected both in principle and in practice.

Setting aside the Bloom et al. (2020) model, the other models estimated by Wang et al. (2022) result in surprisingly similar estimates of elasticities of output with respect to knowledge stocks, and the implied benefit-cost ratios and internal rates of return, though they imply very substantial differences in the structural links between economic outcomes over time and past investments in research—i.e., how this part of the economy actually works. Here, strong priors

[11] Annex Table A-1 includes summary details on the estimates. Annex Figure A-1 displays the alternative lag distribution models.

about the lag relationship based on an informed understanding of the relevant economic aspects of the industries and their institutions—for example, as imposed by Huffman and Evenson (1993) and Alston et al. (2011), and almost all other studies linking agricultural R&D and productivity—are decisive to interpreting the less-definitive and less-discriminating raw statistical evidence.

More than forty years ago, Griliches (1979) suggested:

> "… it is probably best *to assume* a functional form for the lag distribution on the basis of prior knowledge and general considerations and not to expect the data to answer such fine questions. That is, a 'solution' to the multicollinearity problem is a moderation of our demands on the data—our desires have to be kept within the bounds of our means." (Griliches 1979, p. 106, emphasis in original).

That their priors are reasonable (and that those implicit in both the geometric model and the model of Bloom et al. 2020 are not) is evidenced by the detailed documentation provided by Wang et al. (2022) (and in more detail by Pardey et al. 2010 and Alston et al. 2022) regarding (1) the R&D lags (the time spent creating new, relevant agricultural knowledge) in decades-long technology timelines for specific technologies, (2) the decades-long adoption lags between the commercial release of a technology and its maximum economic impact, and (3) the disadoption processes as new machines, varieties, materials and methods progressively replace those in use.

Agricultural MFP and Agricultural R&D Knowledge Stocks

In their Figure 5, Bloom et al. (2020, p. 1120) present a counterpart of their Figure 1, but now using data for U.S. agriculture rather than the entire U.S. economy. Here they seem to have implicitly allowed for a brief (five-year) R&D lag by plotting a smoothed measure of TFP growth over the "next five years" against a measure of current research input equal to nominal total (public and private) U.S. agricultural R&D expenditure divided by the average wage of high-skilled workers. In Figure 2 we plot our own (InSTePP) measures of U.S. total (public and private) agricultural R&D expenditure divided by the InSTePP R&D deflator and U.S. agricultural MFP for the years since WWII (see Alston and Pardey 2022 for further details). Like those provided by Bloom et al. (2020) the plots in Figure 2 show a slowdown in U.S. agricultural productivity growth (documented in detail by Andersen et al. 2019, and Pardey and Alston 2021) against considerable real growth in the annual flow of research spending. However, using lag weights estimated by Wang et al. (2022) the growth path of the agricultural R&D knowledge stock is much more congruous with the growth path of MFP.

[Figure 2. *Productivity Growth and Research Effort in U.S. Agriculture*]

Direct Evidence on Productivity of Agricultural Scientists—U.S. Wheat Breeding

More direct evidence on the productivity of agricultural scientists can be gleaned from a more-focused look at wheat—one of four crops studied by Bloom et al. (2020)—, using some additional data resources available to us. Building on Pardey et al. (1996), Chai et al. (2022) compiled detailed data on the creation and adoption of new wheat varieties in the United States over the past century.[12] These data include annual measures for the period 1919–2019 of area planted to each of a long list of commercially grown wheat varieties in each of 16 wheat producing U.S. states, which together accounted for 89% of U.S. planted wheat area in 2019. For the remaining states for which varietal survey data were not available, we imputed the number of varieties by applying a measure of varietal density (number of "counted" varieties per

[12] Such compilations are not available off the shelf. For the 16 largest wheat growing states, for the period 1919–1984 state-specific data on the area sown to each of the "counted" varieties were taken from quinquennial issues of USDA's *Statistical Bulletins on the Distribution of the Varieties and Classes of Wheat in the United States*. For the period 1985–2019, Chai et al. (2022) also compiled the same data by contacting the statistical offices for each of the respective state departments of agriculture. Their compilation includes a total of 1,353 commercially grown and named varieties whose acreage exceeded 0.1% of each state's planted wheat area in the reporting year, hereafter designated as "counted" varieties.

million planted acres) observed in the 16 main wheat-growing states to the total planted acres for the other states. The resulting measure of the total stock of wheat varieties in use is labeled "Total varieties" in Figure 3a—a tangible measure of genetic "ideas" used in wheat production.[13]

[Figure 3. *Research Effort and Productivity in U.S. Wheat Breeding*]

This measure has generally trended up over the 100-year period, even though planted wheat area has declined steadily in recent decades. While U.S. wheat acreage fell by 48.5% from a peak of 88.3 million acres in 1981 (cf. 77.4 million in 1919) to just 45.5 million acres in 2019, production fell by just 30.6%, from 75.8 to 52.6 million tons. Over the hundred years, 1991–2019, substantial, across-the-board increases in the number, turnover and spatio-temporal diversity of commercially grown wheat varieties enabled yields to grow (linearly) at an average annual rate of 1.03%. As Chai et al. (2022) report, U.S. farmers planted a total of just 33 named varieties in 1919, which increased to 186 varieties in 2019. The corresponding increase in varietal intensity—averaging 0.8 varieties per million acres in 1919 versus 9.1 varieties per million acres in 2019—represents a 10-fold increase in varietal diversity. Improving the match-up between spatially variable growing conditions and locally adapted wheat varieties is key to realizing the genetics-by-environment (G x E) gains that sustain or enhance crop yields.[14]

Production environments (e.g., climate conditions or pest pressures) and market demands (e.g., for specific crop product attributes) are also subject to considerable change, both of which

[13] Chai et al. (2022) use phylogenetically informed approaches to identify "effectively distinct" wheat varieties over the past century, which might be a more defensible counterpart of "ideas" than varieties qualifying for IPRs. Their measure of effectively distinct varieties has also trended up against the fairly constant wheat-breeding effort we identify here.

[14] Referring to an earlier era, Olmstead and Rhode (2002) pointed out that slow or zero yield growth alone could not be construed as evidence of a lack of innovation: simply sustaining U.S. wheat yields as the industry moved many miles north and west, and pests and diseases coevolved, required a continuing stream of ideas and innovations embodied in new varieties.

can act to undermine the value of existing crop varieties. Wheat breeders have enabled growers to address these factors by increasing the diversity (not just the number) of wheat varieties over time. For example, in 1919 just 1.3% of the planted area was sown to new varieties (i.e., ≤ 3 years in use). By 2019, 10.5% of the U.S. wheat area was planted to new varieties, such that the area-weighted age of commercially-grown varieties declined dramatically from 36.4 years in 1919, to 16.0 years in 1960, and down to just 9.3 years in 2019 (Chai et al. 2022). The decline in the area share of older varieties (i.e., >15 years in use) was particularly pronounced; 68.8% in 1919, and only 10.0% in 2019.

Taking a closer look at these data on varieties in use over time, we can see when varieties entered production and when they were supplanted by others. We opted, for present purposes, to define a "new" variety as one that is three years old or younger, such that the total number of varieties aged three years or less, divided by three, is a three-year moving average measure of the number of new varieties per year.[15] In Figure 3a that measure of the flow of new wheat varieties per year fluctuates considerably around a varying trend. The flow of new varieties trends up slightly for three or four decades, from about four per year in the first decade to about twice that number by the mid-1950s. Over the next three decades (say, 1955–1984), it surges from about eight per year in the mid-1950s to three or four times that number by the mid-1980s. Over the next two decades (say, 1985–2004), the rate is fairly flat, at around 20–30 new varieties per year,

[15] The lack of uniformity in the reported statistics makes it difficult to be precise in defining "varietal newness." We record the date of first commercial use, which is not necessarily that same as the date of commercial release (or first availability). Bulking breeder seed to obtain sufficient quantities for plantings at commercial scale takes some time, and there are supply chain frictions in the distribution process. These factors have changed over time and vary over space. With that in mind, we designated a variety as being "new" if it had been reported as commercially grown for up to 3 years.

but then it really surges, to roughly double again over the last 10–15 years depicted, to about 70 per year by the mid-2010s.

Under the Plant Variety Protection Act (PVPA) of 1970, wheat breeders may apply for a plant variety protection (PVP) certificate, which confers limited patent-like protection (see, e.g., Alston and Venner 2002). Hence, beginning in 1971, the total number of new PVP applications is an alternative measure of the flow of new varieties. In Figure 3a it can be seen that the plot of PVP applications has a similar shape to the plot of our measure of new varieties per year, especially in the most recent two or three decades. Each of the two measures reinforces our confidence in the other, as they indicate similar trends.

We have three alternative measures of the corresponding inputs for the years since 1970, derived drawing primarily on data from the USDA Current Research Information System (CRIS), as shown in Figure 3b. One of these input measures, is given by dividing the total annual flow of public and private expenditures on wheat breeding R&D by the InSTePP deflator for R&D spending (shown as "R&D spending" in Figure 3b). Bloom et al. (2020) might refer to this as a measure of the number of researchers, but it is better regarded as an implicit index of the quantity of <u>all</u> R&D inputs, not just the skilled labor input but also other labor, land, capital and materials. Our second measure is a direct measure of the annual input of researcher effort, the full-time-equivalent number of wheat breeding scientists (assistant professor and above) in the public and private sector (shown as "scientific years" in Figure 3b).

A problem with both these measures is that current R&D spending or researcher effort has little or nothing to do with the current release of new varieties, since the R&D process takes years or decades. A better measure of the inputs that contributed to the current release of new wheat varieties is the flows of services from the corresponding R&D knowledge stock, which can be proxied by the stock itself. We computed the "Stock of knowledge" in Figure 3b by applying R&D lag weights from Wang et al. (2022) to CRIS data on annual wheat R&D expenditures projected back to 1935, using the InSTePP data on overall U.S. agricultural R&D spending.

Comparing these three measures, the flow of wheat breeding scientific research effort fluctuates around a fairly flat trend (around 200 scientist-years per year) while the flow of R&D spending fluctuates in similar ways around a somewhat upward trend—from less than $100 million per year (2019 prices) in the early 1970s to around $135 million per year during the later 2010s, after a peak of $189 million in 2012. Reflecting the long lag distribution and much lower annual spending in the years before 1970 than after 1970, the wheat R&D knowledge stock grows more rapidly and more smoothly, from around $22 million in 1970 to $106 million in 2017.

Figure 4 displays six alternative annual measures of wheat research productivity over almost five decades 1970–2018, derived by combining each of the three measures of research inputs from Figure 3b with one or the other of the two measures of research outputs from Figure 3a—either "New varieties" in Figure 4a or "New PVP applications" in Figure 4b. In each case the best measure uses the R&D stock in the denominator. In Figure 4a, the rate of new varieties per dollar of R&D stock fluctuates around a declining trend until the early 1990s, followed by a flat trend until about 2000, after which it trends up. In Figure 4b, the rate of new PVPs per dollar of R&D stock fluctuates around a flat trend until the early 1990s after which it trends up. Using either measure of research output, productivity of wheat research has not been falling over the past 20–30 years and, in more recent years, if anything, it has been increasing. Moreover, if we had used either of the other two (less favored) measures of research inputs, flows of real spending per year or number of scientific-years per year, the evidence against a decline in wheat research productivity is even stronger.

[Figure 4. *Wheat Research Productivity, Various Measures, 1970–2017*]

Meta-Evidence on Rates of Return to Research

An alternative measure of research productivity is the rate of return to research. Rao et al. (2019) report the results from a meta-analysis encompassing 492 studies published since 1958 that collectively reported 3,426 estimates of rates of return to agricultural R&D. They conclude that "… the contemporary returns to agricultural R&D investments appear as high as ever" (Rao et al. 2019, p. 37); see also Alston, Pardey and Rao (2021). In general, improvements in the technology of science and in the human capital of scientific researchers have made researchers more productive, and it seems these gains in research productivity have been sufficient to offset any decline caused by other factors. These recent findings echo results from various previous reports regarding the links between research and productivity—whether in agriculture and the broader economy—and whether agricultural or other industrial research was becoming less productive in this sense (see, e.g., Griliches 1979, 1992, 1996).

III. Conclusion

Bloom et al. (2020) assume research-induced productivity gains last unabated forever—i.e., an infinite lag. In the agricultural economics literature, this special case—with an infinite overall lag length and an instantaneous research impact—is strongly rejected by the evidence from narrow studies of particular technologies (e.g., such as wheat varieties) as well as broad studies of national agricultural research systems and anything in-between. The same must be true in applications outside agriculture. In brief, there is no empirical support in the existing literature for the R&D lag distribution model implicitly assumed by Bloom et al. (2020), either in terms of the realism of the assumptions (research and adoption really do take time, and at least some technologies do become obsolete) or in empirical evidence from studies that estimated the lag structures, whether in agriculture or other parts of the economy.

Using more defensible models of the time-consuming process of knowledge creation and adoption (and eventual disadoption) of the resulting innovations, applied to detailed data for U.S. agriculture, we see no evidence of a persistent, long-run decline in research productivity, and no evidence that ideas are getting harder to find in that sense. No doubt, innovations in the technology of agricultural science have contributed considerably to sustaining the productivity of agricultural scientists.

Crop improvement especially has benefited from research productivity-enhancing innovations such as marker-assisted breeding, improved gene manipulation technologies (e.g., genetic engineering or gene editing), and enhanced quantitative genetics methods using ever-larger high-performance-computing facilities. And, in the case of wheat research, constant or increasing research productivity combined with an increasing wheat-breeding research effort has resulted in an acceleration in the annual flow of new varieties, and an increase in the productivity of wheat breeders by this measure—at least since 1990. Even so, wheat yields have grown only linearly such that proportional yield growth has slowed. But as we have shown, we cannot interpret the juxtaposition of slowing yield growth with increased research effort directly as evidence that it is getting harder to create new varieties. It is much harder than that to know—or, indeed, to know how to determine—if ideas are getting harder to find. But an important first step is to have a reasonably defensible conception of the knowledge creation process before we set out to check for retrograde changes in the structure of that process.

IV. References

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–51.

Alston, J.M., M.A. Andersen, J.S. James and P.G. Pardey. *Persistence Pays: U.S. Agricultural Productivity Growth and the Benefits from Public R&D Spending.* New York: Springer, 2010.

Alston, Julian M., Matthew A. Andersen, Jennifer S. James and Philip G. Pardey. 2011. "The Economic Returns to U.S. Public Agricultural Research." *American Journal of Agricultural Economics* 93 (5): 1257–77.

Alston, Julian M. and Philip G. Pardey. 2021. "The Economics of Agricultural Innovation." In *Handbook of Agricultural Economics*, Volume 5 edited by Christopher B. Barrett and David R. Just, 3895–3980. Amsterdam: Elsevier.

Alston, Julian M., Philip G. Pardey, Devin Serfas and Shanchao Wang. 2022. "Research Lags Redux: Reconciling Industrial and Agricultural R&D Lag Models." InSTePP Working Paper. St Paul, MN: International Science and Technology Practice and Policy Center, University of Minnesota (in process).

Alston, Julian M., Philip G. Pardey, and Xudong Rao. 2021. "Payoffs to a Half Century of CGIAR Research." *American Journal of Agricultural Economics* published online https://doi.org/10.1111/ajae.12255

Alston, Julian M., and Raymond J. Venner. 2002. "The Effects of the U.S. Plant Variety Protection Act on Wheat Genetic Improvement." *Research Policy* 31(4): 527–542.

Andersen, M.A., J.M. Alston, P.G. Pardey and A. Smith. "A Century of U.S. Farm Productivity Growth: A Surge Then a Slowdown." *American Journal of Agricultural Economics* 100(4)(2018): 1072-1090.

Bureau of Economic Analysis (BEA). 2022. "National Income and Product Accounts." Federal Reserve Economic Data (Fred). https://fred.stlouisfed.org/series/Y006RC1Q027SBEA and https://fred.stlouisfed.org/series/Y057RC1A027NBEA (accessed April 3, 2022)

Bureau of Economic Analysis (BEA). 2017. "National Income and Product Accounts." Federal Reserve Economic Data (Fred). https://fred.stlouisfed.org/ (accessed January 31, 2017).

Bureau of Labor Statistics (BLS). 2017. "Historical Multifactor Productivity Measures, 1948-2014." https://www.bls.gov/mfp/home.htm (accessed February 17, 2017)

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–44.

Chai, Yuan, Philip G. Pardey and Kevin A.T. Silverstein. 2022. "Scientific Selection: A Century of Increasing Crop Varietal Diversity in U.S. Wheat." InSTePP Working Paper. St Paul, MN: International Science and Technology Practice and Policy Center, University of Minnesota (in process).

Evenson, Robert E. 1967. "The Contribution of Agricultural Research to Production." *Journal of Farm Economics* 49(5): 1415–1425.

FAO (Food and Agriculture Organization of the United Nations). 2021. FAOSTAT Statistical Database. Value of Agricultural Production. Rome, Italy: FAO. Downloaded from https://www.fao.org/faostat/en/#data/QV on December 21, 2021.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth*. Princeton: Princeton University Press.

Griliches, Zvi. 1979. "Issues in Assessing the Contribution of R&D to Productivity Growth." *Bell Journal of Economics* 10(1): 92-116.

Griliches, Zvi. 1992. "The Search for R&D Spillovers." *Scandinavian Journal of Economics* 94 (Supplement): 29–47.

Griliches, Zvi. 1996. *R&D and Productivity: The Econometric Evidence*. Chicago: Chicago University Press.

Grossman, Gene M., and Elhanan Helpman. 1991. *Innovation and Growth in the Global Economy.* Cambridge, MA: MIT Press.

Kortum, Samuel. 1997. "Research, Patenting, and Technological Change." *Econometrica* 65 (6):1389–1419.

Hall, Bronwyn H., Jacques Mairesse, and Pierre Mohnen. 2010. "Measuring the Returns to R&D." In *Handbook of the Economics of Innovation*, Vol. 2, edited by Bronwyn H. Hall and Nathan Rosenberg, 1033–82. Amsterdam: Elsevier.

Huffman, Wallace E., and Robert E. Evenson. 1993. *Science for Agriculture: A Long-Term Perspective*. Ames: Iowa State University Press.

InSTePP (International Science and Technology Practice and Policy) Center. 2020. Unpublished InSTePP R&D deflator. St Paul, MN: University of Minnesota.

InSTePP (International Science and Technology Practice and Policy) Center. 2022. Unpublished InSTePP wheat varietal database. St Paul, MN: University of Minnesota.

Jones, Benjamin F. and Lawrence H. Summers. 2020. "A Calculation of the Social Returns to Innovation." NBER Working Paper 27863. Cambridge, MA: National Bureau of Economic Research.

Jones, Charles I. 1995. "R&D Based Models of Economic Growth." *Journal of Political Economy* 103 (4): 759–84.

Jones, Charles I. 2005. "Growth and Ideas." In *Handbook of Economic Growth,* Vol. 1B, edited by Philippe Aghion and Steven A. Durlauf, 1063–1011, Oxford: North Holland.

Li, Wendy C.Y., and Bronwyn H. Hall. 2020. "Depreciation of Business R&D Capital." *Review of Income and Wealth* 66 (1):161-180. doi: 10.1111/roiw.12380 (also NBER Working Paper 22473).

Olmstead, Alan L. and Paul W. Rhode. 2002. "The Red Queen and the Hard Reds: Productivity Growth in American Wheat, 1800–1940." *Journal of Economic History* 62 (4): 929–66.

Pardey, Philip G., Julian M. Alston, Jason E. Christian and Shenggen Fan. *Hidden Harvest: U.S. Benefits from International Research Aid*, Food Policy Statement No. 23, Washington, D.C.: International Food Policy Research Institute, September 1996.

Pardey, Philip G. and Julian M. Alston. 2021. "Unpacking the Agricultural Black Box: The Rise and Fall of American Farm Productivity Growth." *The Journal of Economic History* 81 (1): 114–155. doi:10.1017/S0022050720000649

Pardey, Philip G., Julian M. Alston and Vernon W. Ruttan. 2010. "The Economics of Innovation and Technical Change in Agriculture." Chapter 22 in B.H. Hall and N. Rosenberg, eds. *Handbook of the Economics of Innovation*, Amsterdam: Elsevier.

Rao, Xudong, Terrance M. Hurley, and Philip G. Pardey. 2019. "Are Agricultural R&D Returns Declining and Development Dependent?" *World Development* 122: 27–37. Available at https://doi.org/10.1016/j.worlddev.2019.05.009

Romer, Paul M. 1990. "Endogenous Technical Change." *Journal of Political Economy* 98 (5): S71–102.

Serfas, Devin, Julian M. Alston, and Philip G. Pardey. 2022. "The Returns to Industrial R&D: A Comprehensive Meta-Assessment and Partial Critique of the Evidence." InSTePP Working Paper. St Paul, MN: International Science and Technology Practice and Policy Center, University of Minnesota (in process).

USDA-AMS (Agricultural Marketing Service). 2022. PVPO Application Status. US PVP Application Status Report. Washington, DC: United States Department of Agriculture. Downloaded on February 4, 2022 from https://www.ams.usda.gov/services/plant-variety-protection/application-status

USDA-ERS (Economic Research Service). 2022. Agricultural Productivity in the United States. Washington, DC: United States Department of Agriculture. National Tables, Updated 1/6/2022. Downloaded on April 9, 2022 from https://www.ers.usda.gov/data-products/agricultural-productivity-in-the-u-s/agricultural-productivity-in-the-u-s/#National%20Tables

USDA-CRIS (Current Research Information System). 2020. Unpublished CRIS data files, authors extraction. Washington, DC: United States Department of Agriculture, Current Research Information System.

Wang, Shanchao, Julian M. Alston, and Philip G. Pardey. 2022. "R&D Lags in Economic Models: Theory and Assessment using Data for U.S. Agriculture." InSTePP Working Paper. St Paul, MN: International Science and Technology Practice and Policy Center, University of Minnesota.

Figure 1. *U.S. Post-WWII MFP Growth, Research Effort, and Research Productivity*

a. MFP Growth and Research Effort

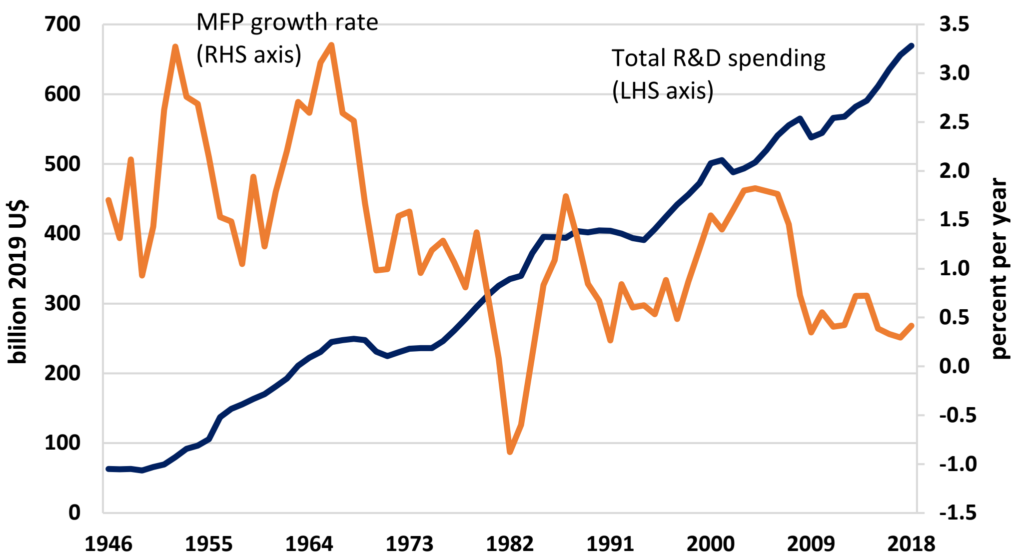

b. Research Productivity (MFP Growth per Research Effort)

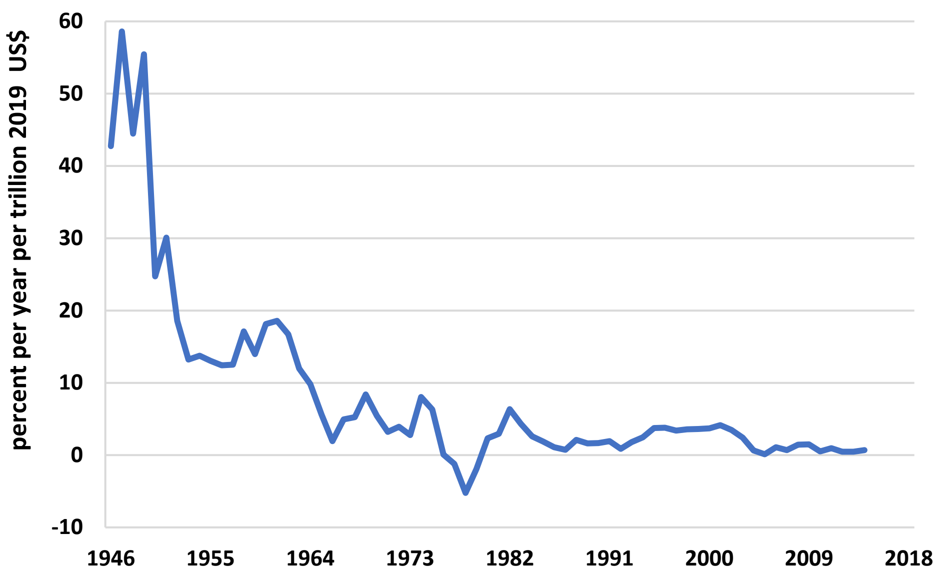

Sources: Authors' calculations based on private business MFP data series from Pardey and Alston (2021) for MFP growth rates. Total R&D spending from 1929 to 2018 is from BEA (2022). We extended the BEA series to 1890 using predicted ratios of total AgR&D in total R&D.

Notes: Panel a corresponds to Bloom et al. (2020, Figure 1) and Panel b corresponds to Bloom et al. (2020, Figure 2). MFP growth rate is a five-year forward moving average of the log difference year to year growth rate of U.S. private business MFP from Pardey and Alston (2021). R&D spending is a quantity measure of research effort, equal to total R&D spending divided by the InSTePP R&D deflator (based = 1.0 in 2019).

Figure 2. *Post-WWII MFP Growth and Research Effort in U.S. Agriculture*

a. AgMFP Growth and Ag Research Effort

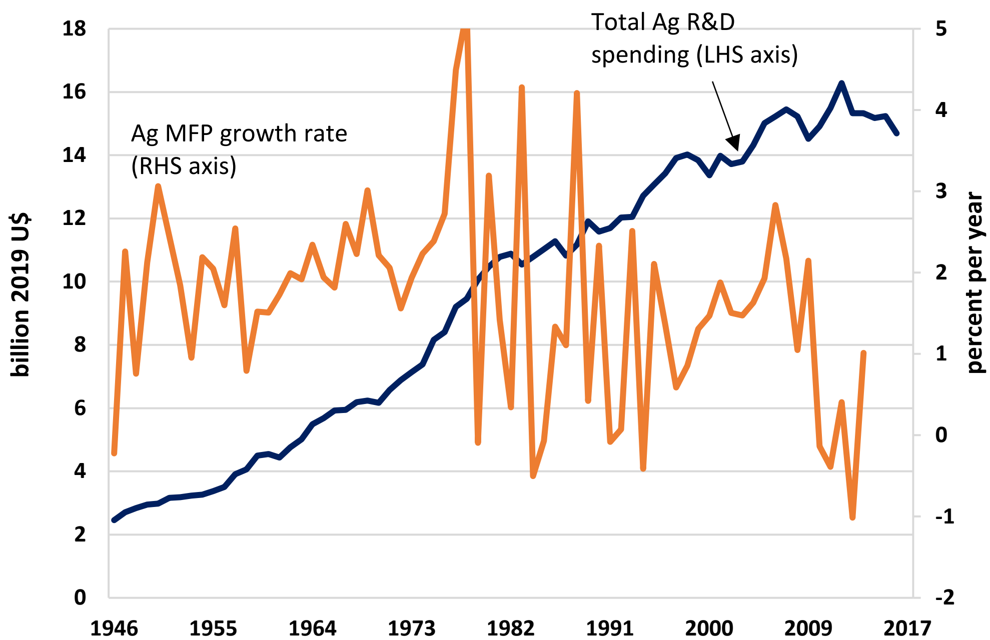

b. Research Productivity (AgMFP Growth per Ag Research Effort)

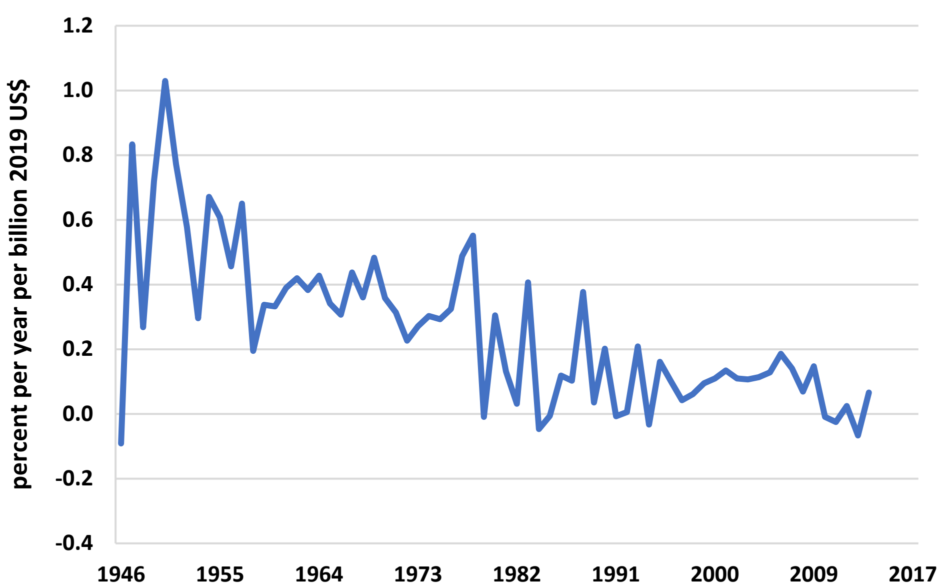

Sources: Authors' calculations based on agricultural MFP data series from Pardey and Alston (2021) and USDA-ERS (2022) for AgMFP growth rates. Total AgR&D spending is from InSTePP data on U.S. agricultural R&D spending deflated by the InSTePP R&D deflator (based = 1.0 in 2019).

Notes: Panel a: The rate of change in the national ag TFP series from USDA-ERS (2022) was used to project the InSTePP agMFP series forward from 2008 to 2019. Panel b: 5-year forward moving average of AgMFP per real 2019 AgR&D spending

Figure 3. Research Effort and Productivity in U.S. Wheat Breeding

a. A Century of U.S. Wheat Varieties (Ideas), 1919–2019

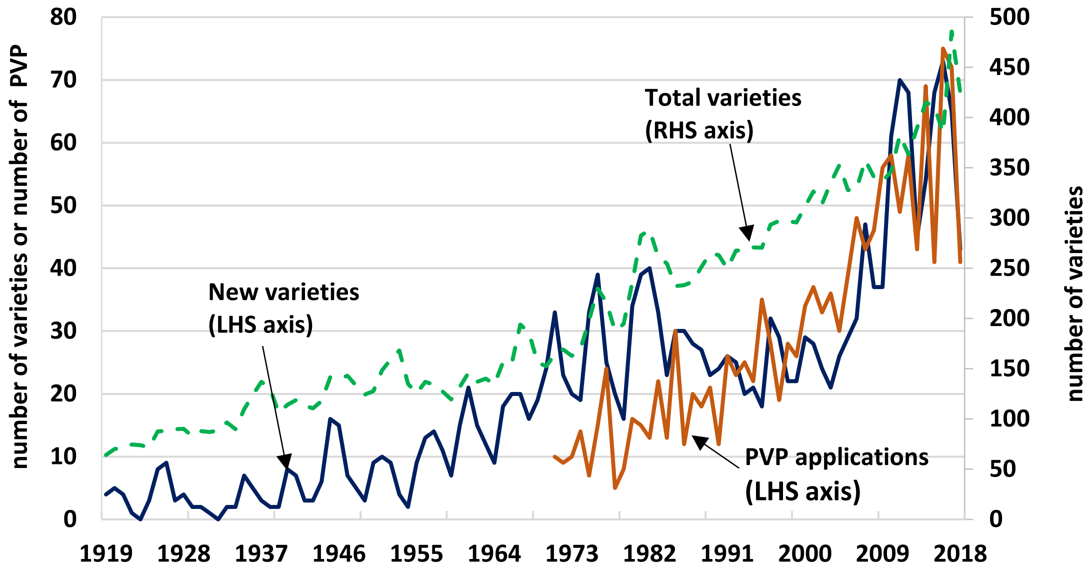

b. Inputs to the Production of New U.S. Wheat Varieties, 1970–2017

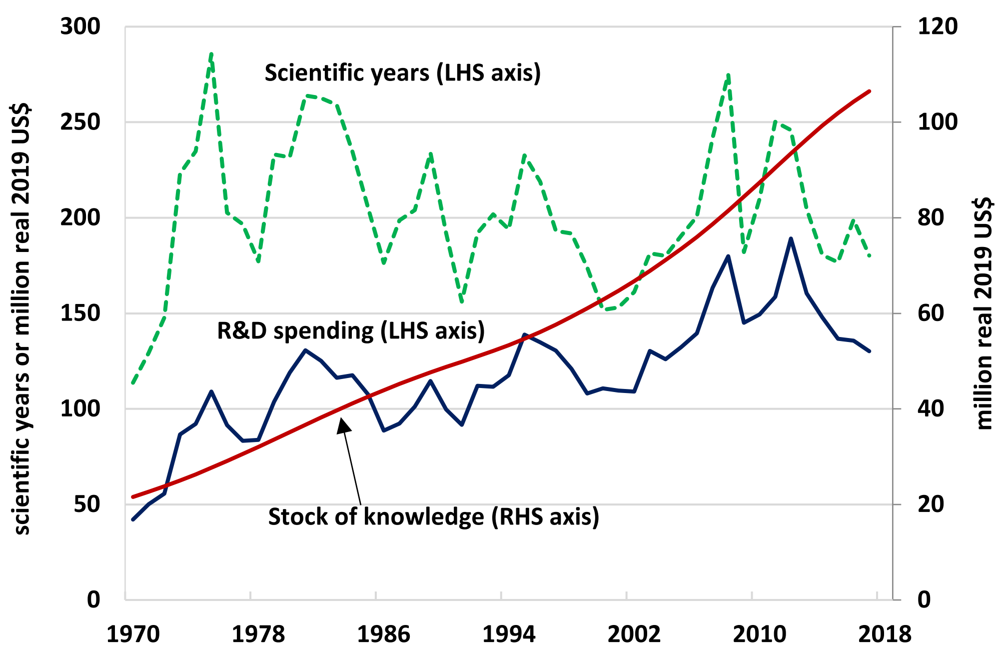

Sources: Authors' calculations based on the InSTePP wheat varietal database (Chai et al. 2022) for the number of varieties; on USDA Agricultural Marketing Service (2022) for PVP applications (PVPs); and on USDA-CRIS unpublished data files for R&D; R&D deflator is from InSTePP (2020).

Notes: "New varieties" represent the total number of "counted" varieties, in use for 3 years or less, planted in 16 wheat-producing U.S. states for which data are available between 1919 and 2019. These 16 states accounted for 89% of planted wheat area in 2019. Total varieties is the total number of "counted" varieties planted in the United

States. We applied the variety density (number of varieties per million acres) for the 16 states to the total planted acres in the remaining wheat-growing states to estimate their counts of varieties. "PVP applications" is the number of applications for PVP certificates for wheat varieties by U.S. applicants.

Scientific years is the number of years of scientist (assistant professor and above) effort in the public and private sectors. The number of wheat-breeding scientific years for the public sector was extracted from USDA-CRIS unpublished data files. To estimate wheat breeding scientific years for the private sector, we divided private wheat-breeding R&D spending by the ratio of public wheat-breeding R&D spending to public scientific years. Thus, we assumed that the ratio of R&D spending-to-FTE is the same in the public and private sectors.

Private wheat-breeding R&D expenditures were estimated by applying the average 1970–1990 share of public wheat-breeding in total public wheat research to all wheat private R&D spending. For the public sector, the wheat R&D spending data and wheat-breeding R&D data were extracted from USDA-CRIS unpublished data files. Total wheat private R&D spending was calculated using the following formula: total wheat private R&D spending = 0.9323 x total private ag R&D x (wheat VOP/ag VOP). Data series on U.S. wheat VOP (value of production) and ag VOP were taken from FAOSTAT (2021). Total wheat-breeding R&D spending is the sum of wheat-breeding R&D spending by the public and the private sectors, deflated by the InSTePP R&D deflator (=1 in 2019).

The wheat-breeding R&D stock was estimated by applying lag weights from the 50-year gamma lag distribution model estimated by Wang et al. (2022, model 1) using U.S. national data. This lag distribution model has an overall shape very similar to that of the trapezoidal model, and a shorter mean lag than the gamma lag model from Alston et al. (2010, 2011), estimated using U.S. state-level data in a slightly different specification and including fewer years of data. See Annex for further details.

Figure 4. *Wheat Research Productivity, Various Measures, 1970–2017*

a. New varieties

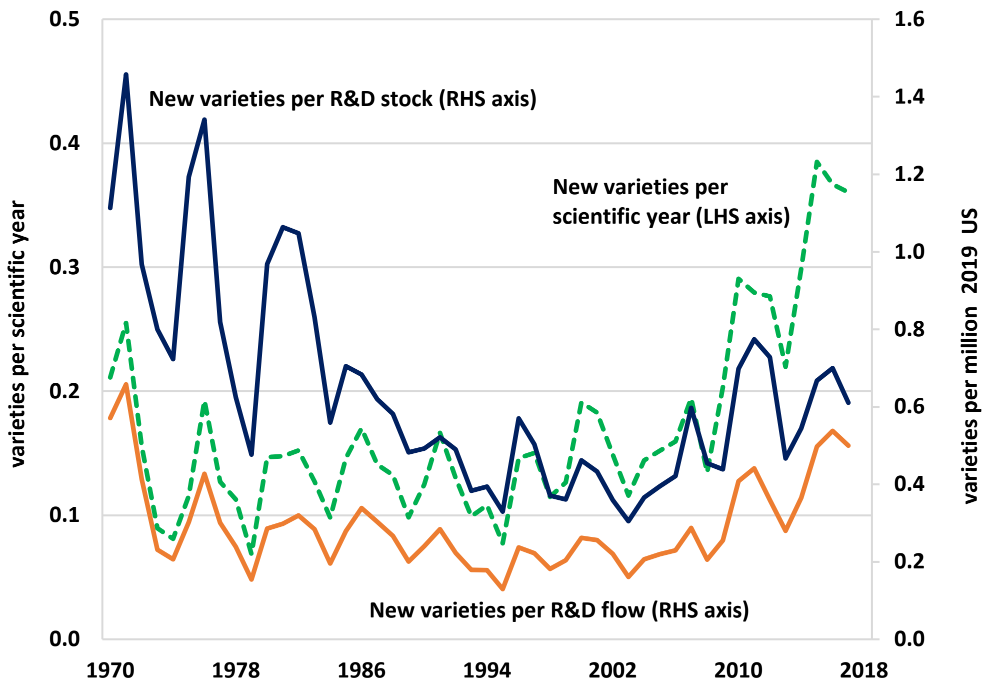

b. New PVP applications

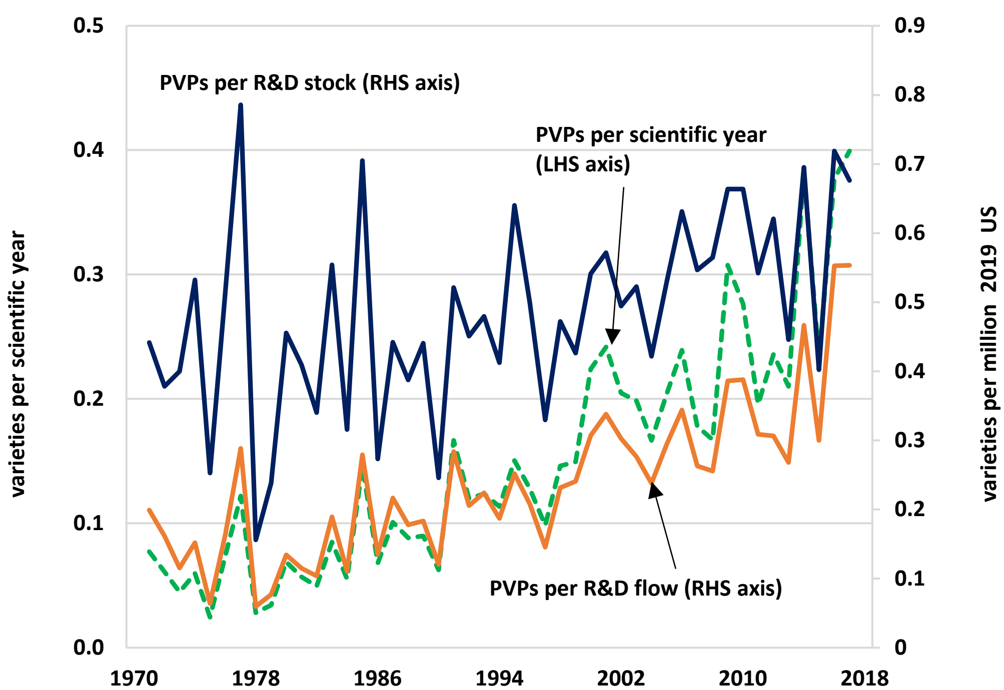

Sources: See Figure 3.

Notes: New varieties is the number of counted wheat varieties ≤ 3 years in use, planted in the principal U.S. wheat-producing states (Figure 3a). PVPs is the total number of applications for PVP certificates for public and private wheat varieties by U.S. applicants (Figure 3a). Scientist years represents the total of public and private wheat-breeding scientist years (Figure 3b). R&D flow is public and private wheat-breeding R&D deflated by the InSTePP R&D deflator (Figure 3b). R&D stock is the total public and private wheat-breeding knowledge stock (Figure 3b). See notes under Figure 3 for further details.

Annex: Supporting Material

Are Ideas *Really* Getting Harder to Find?

Julian M. Alston and Philip G. Pardey

Figure A-1. *R&D Lag Distribution Models Reported by Wang et al. (2022)*

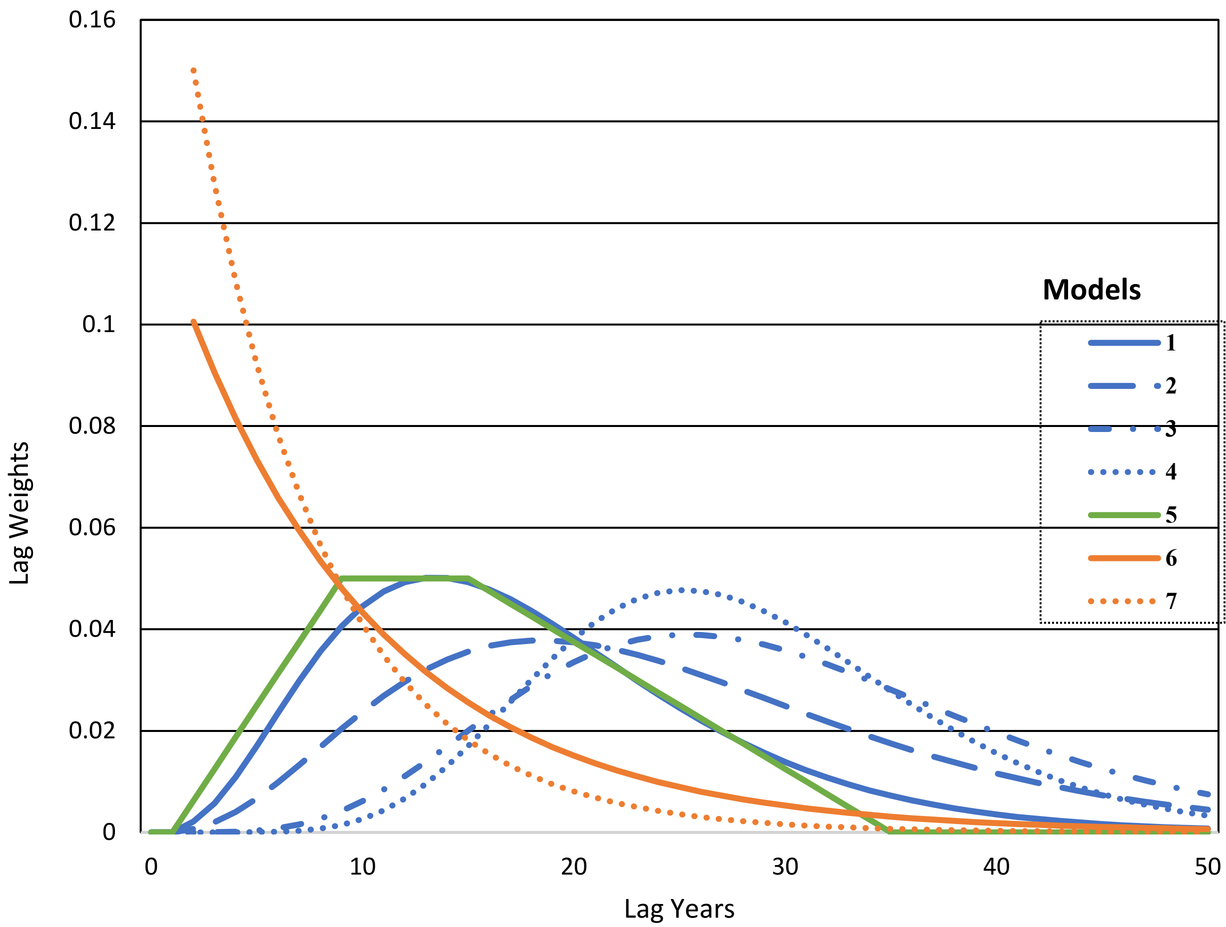

Sources: Developed by the authors.

Notes: Gamma lag distribution models (Models 1–4) are shown in blue; the trapezoidal lag distribution (Model 5) is shown in green; geometric lag distribution models (Models 6 and 7) are shown in orange; the Romer-Bloom model (Model 8) is not depicted here. See Table A-1 for a summary of the parametrizations of these lag distribution models and Table 4 of Wang et al. (2022) for more complete details.

Table A-1: *Models of U.S. Agricultural MFP using Various R&D Lag Models to Represent Knowledge Stocks*

		Regressors						*Time-Series Tests*
Model	*Lag model (parameters)*	*Constant* (1)	$\ln(K_t)$ (2)	K_t (3)	W_t (4)	T_t (5)	SSE	$\ln(K_t)$ is I(1) (Coint. Tests)
1	Gamma	2.772***	0.307***		0.026***	0.008***	0.074	Pass
	(0.75, 0.80)	(0.550)	(0.084)		(0.004)	(0.003)		(Pass)
2	Gamma	2.873***	0.304***		0.026***	0.008**	0.083	Pass
	(0.75, 0.85)	(0.577)	(0.092)		(0.004)	(0.003)		(Pass)
3	Gamma	3.304***	0.247**		0.026***	0.009**	0.096	Pass
	(0.85, 0.80)	(0.559)	(0.094)		(0.004)	(0.004)		(Pass)
4	Gamma	3.373***	0.235**		0.026***	0.009**	0.098	Fail
	(0.90, 0.70)	(0.586)	(0.098)		(0.004)	(0.004)		(Pass)
5	Trapezoidal	2.798***	0.299***		0.026***	0.009***	0.073	Fail
		(0.561)	(0.084)		(0.004)	(0.003)		(Fail)
6	Geometric	3.222***	0.227**		0.026***	0.013***	0.077	Pass
	($\delta = 0.10$)	(0.603)	(0.087)		(0.004)	(0.002)		(Fail)
7	Geometric	3.348***	0.205**		0.026***	0.013***	0.078	Pass
	($\delta = 0.15$)	(0.592)	(0.084)		(0.004)	(0.002)		(Fail)
8	Romer–Bloom	4.781***		-1.02E-06	0.026***	0.023***	0.100	Fail
		(0.018)		(5.26E-07)	(0.004)	(0.002)		(Fail)
				[-0.115]				

Sources: Wang et al. (2022, Table 6).

Notes: Results were obtained using the Prais-Winsten Procedure to correct for autocorrelation. Eicker-Huber-White standard errors in parentheses in columns (1) through (5). ***, **, and * denote significance levels of 0.01, 0.05 and 0.10, respectively. Coefficients in column (3) are elasticities of MFP with respect to the knowledge stock. In the Romer-Bloom model the elasticity is shown in square brackets in column (4), calculated at the median of constructed Romer-Bloom knowledge stock across the period 1940–2007. The gamma model parameters (γ, λ) are used to compute the lag weights: $b_k = \frac{(k-g+1)^{\gamma/\{(1-\gamma)\}}\lambda^{(k-g)}}{\sum_{k=g+1}^{50}[(k-g+1)^{\gamma/(1-\gamma)}\lambda^{(k-g)}]}$ for $g < k \leq 50$; otherwise $b_k = 0$.