The Millennium Droughts and Australian Agricultural Productivity Performance: A Nonparametric Analysis

Robert G. Chambers, Simone Pieralli, and Yu Sheng

With the turn of the century, Australian agricultural productivity growth slowed dramatically. We investigate the connection between this slowdown and climatic factors by comparing regional-level growth patterns before and after the advent of the Australian Millennium Droughts. The analysis incorporates climatic variates directly into the productivity accounting framework to reflect the stochastic nature of agricultural production, and measured productivity growth is decomposed into four components: technological change, weather-related change, input-scale adjustment, and diffusion (adaptation). Nonparametric productivity measurement and statistical techniques are used to quantify and examine the patterns of the observed productivity slowdown. The analysis suggests that the primary determinant of the slowdown is not a slowdown in technological innovation but climatic-related changes in the pattern and rate of diffusion of technological advances.

Key words: Agricultural TFP, Australian broadacre agriculture, Millennium Droughts, technological innovation and diffusion.

JEL codes: D24, Q11, Q54.

Strong agricultural productivity growth in the OECD countries has proven a key driver of global development. A growing body of evidence suggests, however, that agricultural productivity growth has slowed dramatically since 2000 (Thirtle et al. 2004; Alston, Babcock, and Pardey 2010; Sheng, Mullen, and Zhao 2010; Alston, Andersen, and Pardey 2015 and Sheng et al. 2020). For example, Sheng et al. (2020) report that OECD agricultural productivity grew at a sluggish annual rate of 0.5% after 2000, or approximately one-third of its average since 1973.

Australia seems a prime example. Between 1979 and 2013, its agricultural productivity grew at an annual average rate of 1.6%, putting it only slightly behind US agriculture, which grew at an average rate of 1.69%. That strong performance, however, masks a severe slump. Between 1979 and 1994, the annual growth rate was 2.8%, while after that it was 0.6% (Sheng, Yang, and Zhao 2018).

Popular explanations for this pronounced slump include climatic conditions. Particular emphasis has been given to the *Millennium Droughts* (Stern 2007; Garnaut 2011; BOM 2015; Hughes, Lawson, and Valle 2017). For example, Hughes, Lawson, and Valle (2017) estimated that drought-related conditions decreased agricultural output by 30% between 2003 and 2011 (BOM 2015; Hughes, Lawson, and Valle 2017). These findings accord with other studies linking poor agricultural performance and weather-related phenomena (Mendelsohn, Nordhaus, and Shaw 1994; Schlenker, Hanemann, and Fisher 2005; Deschênes and Greenstone 2007; Fisher et al. 2012; Cárdenas Rodríguez, Haščič, and Souchier 2016; Yang and Shumway 2016; Liang et al. 2017; Ortiz-Bobea, Knippenberg, and Chambers 2018).

The simultaneous occurrence of the Millennium Droughts and the Australian productivity slowdown defines a natural experiment that facilitates investigating the interaction between

Robert G. Chambers is a professor of Agricultural and Resource Economics at the University of Maryland. Simone Pieralli is a senior lecturer in Farm Management in the Agriculture Group, School of Agriculture and Environment, Massey University, New Zealand. Yu Sheng is an associate professor and deputy director of China Centre for Agricultural Policy, School of Advanced Agricultural Sciences, Peking University and Senior Economist, Australian Bureau of Agricultural and Resource Economics and Sciences. We acknowledge financial support from the National Natural Sciences Foundation of China (71873005), PKU-NSFC (2018KEY06), ABARES (ATMAC 43643), the Data Centre for Management Science Peking University, and the Bill & Melinda Gates Foundation (OPP1193805).
Correspondence to be sent to: yu.sheng@pku.edu.cn; ysheng.ccap@pku.edu.cn

Amer. J. Agr. Econ. 102(5): 1383–1403; doi:10.1111/ajae.12098
Published online April 28, 2020

extreme climatic conditions and productivity performance. Using techniques developed by Chambers and Pieralli (2020), we exploit that natural experiment by decomposing measured annual productivity growth into four components (technological change, weather-related effects, input/scale effects, and technological diffusion) and analyzing how their performance has varied with advent of the Millennium Droughts. As in Chambers and Pieralli (2020), the analysis recognizes agricultural production's stochastic nature by incorporating measured weather variates directly into the productivity accounting. Unlike Chambers and Pieralli (2020), who study long-term changes (over three decades) in productivity performance, we compare the behavior of year-to-year productivity growth for a pre-Millennium-Droughts period and a Millennium-Droughts period. That allows us to examine how the advent of the extreme weather conditions associated with the Millennium Droughts affected short-term patterns of agricultural productivity growth. To ensure robustness to varying assumptions about functional structure and the underlying data-generating processes, the analysis relies on nonparametric productivity measurement techniques and nonparametric statistical analyses. The empirical analysis uses data on broadacre Australian agriculture for thirty-two production regions contained in three production zones for 1979–2013. Our results suggest that the observed productivity slowdown is not statistically associated with a slowdown in the average rate of technological change. Instead, weather-induced changes in patterns of technological diffusion, which differ across the three production zones, seem to have played a more prominent role.

The paper proceeds as follows. The next section discusses historical Australian productivity performance. We then develop the analytic model, the productivity decomposition, and the empirical method for approximating the aggregate technology for Australian non-irrigated agriculture. A description of the data follows, and then we present our empirical results. The paper then closes.

Australian Non-Irrigated Agricultural Productivity Performance

Non-irrigated (broadacre) agriculture dominates Australian agriculture. In 2013, it generated output worth A$48.5 billion, or approximately 70% of total agricultural output. Broadacre agriculture consists of 128–129 thousand farms distributed over thirty-two regions in three climatic zones (High-Rainfall, Wheat-Sheep, and Pastoral; ABARES 2018, see figure 1). Its non-irrigated nature ensures that output performance is stochastic and depends on natural occuring inputs, such as soil moisture and temperature, that are beyond the direct control of producers.

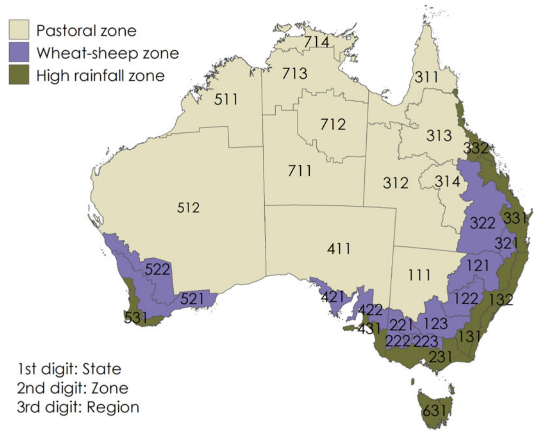

Figure 1. Non-irrigated (Broadacre) agriculture regions in Australia

Source: ABARES (2018).

Figure 2(a) depicts Australian-broadacre total factor productivity (TFP) behavior between 1978 and 2014. TFP increased over that period from a base of 1 in 1978 to 1.47 in 2013. But as figure 2(a) also shows, the pattern of that growth changed dramatically over time. After 1994, productivity growth slowed dramatically. Where it had grown at an annual average rate of 2.8% before 1994, after it only grew at 0.6%.

In retrospect, therefore, the extreme drought conditions experienced in 1994 seem to have marked the beginning of an era of increased drought frequency and slow average productivity growth. Prior to 1994, Australia had experienced two periods of particularly severe and prolonged droughts. These were the "Federation Drought" of 1895–1903 and the "Forties Drought." Repeated instances of extreme droughts followed the 1994 drought, and these increasingly frequent droughts became known collectively after 2002 as the "Millennium Droughts" (Productivity Commission 2009; Sheng and Xu 2019).[1]

[1] The 2003 drought was originally coined the *Millennium Drought*. Henceforth, we use the plural form to denote the successive droughts that have occurred since then. The Productivity Commission's (2009) *Government Drought Support* report contains an extensive discussion of the instances of prolonged drought in Australia.

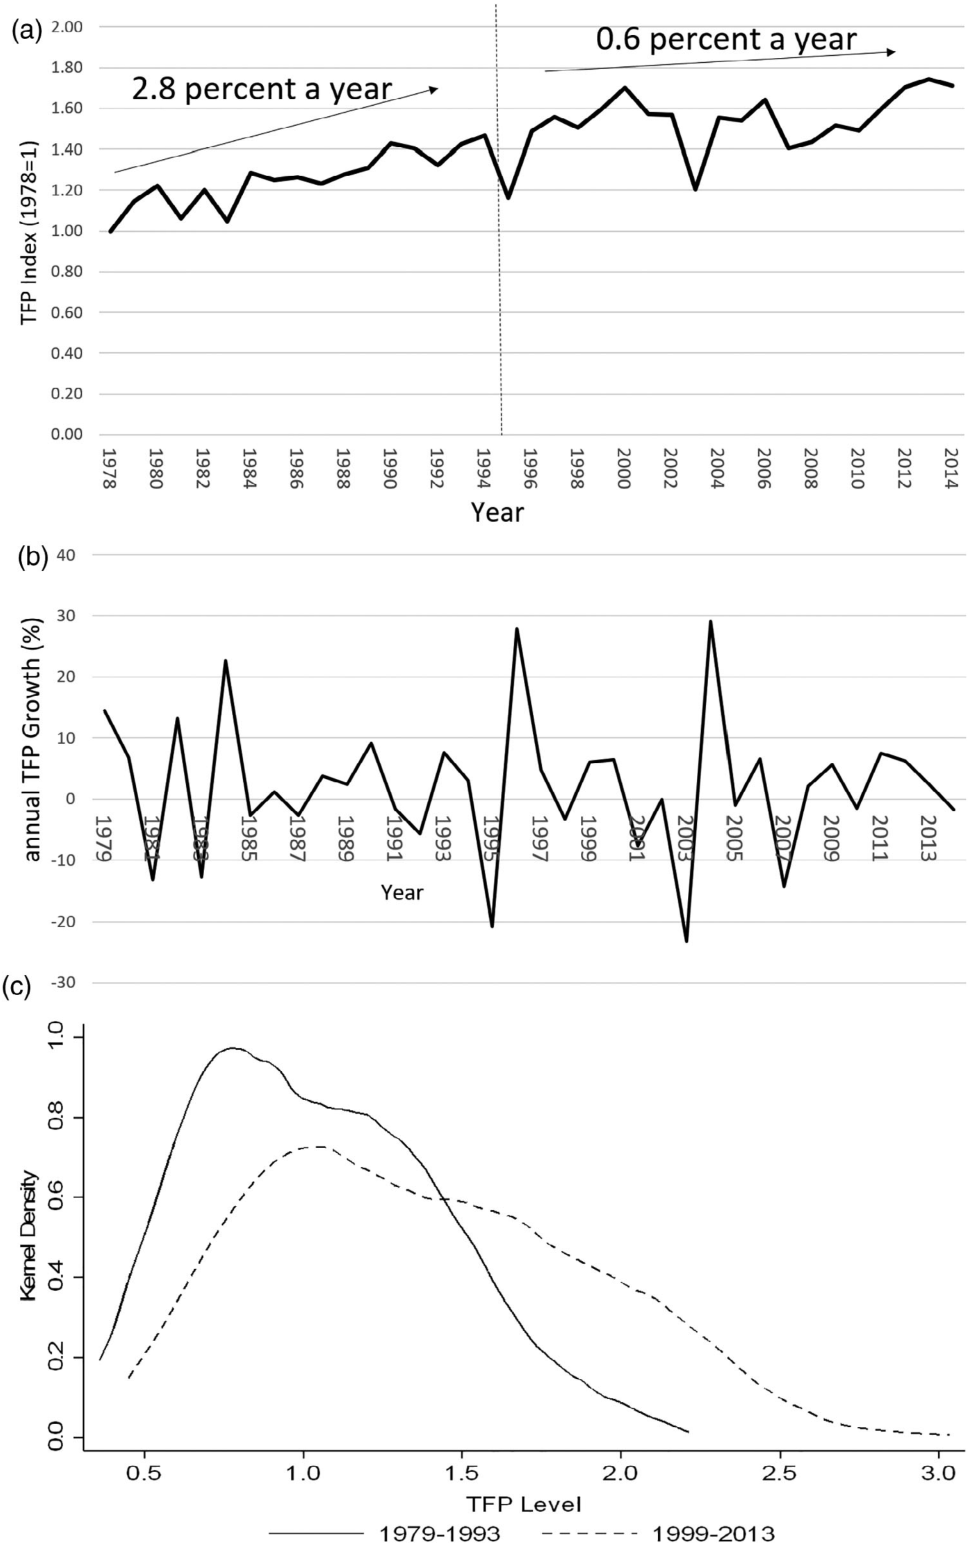

Figure 2. Annual TFP index and growth rate in 1978–2014 and smoothed kernel densities of Australian TFP in 1979–93 and 1999–2013. (a) Australian broadacre agricultural TFP index: 1978–2014 (1978 = 1). (b) Annual TFP growth in Australian broadacre agriculture: 1979–2014 (%). (c) Smoothed kernel densities of TFP in 1979–93 and 1999–2013 (1978 region 111 = 1.0)

Figure 2(b) shows that after 1994 agricultural productivity growth slowed and became more variable. Oscillatory annual percentage changes in productivity (in absolute value terms) exceeded 5% in twelve out of eighteen years. If, as is traditional, measured TFP change is identified with technological change, this observed behavior signals a rapid oscillation between periods of technical progress and technical regress. Literally, this would require the average farmer to make major technical advances only to forget them a year or so later. Absent an assumption of periodic mass psychosis, this seems implausible. Instead, one naturally suspects that other factors, likely beyond producer control and not factored into current productivity accounting, play a significant role in this observed behavior. Weather is a prime suspect.

Another perspective on this volatility is given by comparing productivity performance across the thirty-two regions in three climatic zones before 1994 and after 1999. Figure 2(c) depicts smoothed kernel density estimates of the distribution of regional agricultural productivity performance for 1979–93 and 1999–2013. Comparing the two distributions visually, one sees that the latter has shifted to the right relative to the former. That reflects an increase in average productivity performance. But the 1999–2013 distribution is also more platykurtic than the 1979–93 distribution. Where the 1979–93 distribution contains a hint of bimodality, none appears in the latter. Instead, less mass is concentrated near the mean of the latter distribution and more in its tails.

The Model

Productivity accounting is conducted at an aggregate level. It traditionally assumes that an aggregate production function that relates aggregate output to aggregate input exists. We maintain that assumption. Our method departs from traditional productivity analysis by including weather variates directly into that production function. This is done to recognize the stochastic nature of broadacre agricultural production. Thus, the aggregate production function for broadacre agriculture is denoted as $f(x, w)$, where x is the aggregate input and w refers to a two-vector of weather variates (a soil-moisture and a temperature variate in this study).

The standard productivity measure is aggregate observed output, denoted by y, divided by the aggregate input x. Instead of measuring productivity levels, however, productivity accounting techniques measure period-to-period productivity changes to create indexes relative to a base period. Thus, if there are two periods t and 0, the productivity index in period t relative to that in period 0 (the base period) is given by

$$(1)\qquad TFP(t,0) = \left(\frac{y_t}{x_t}\right) / \left(\frac{y_0}{x_0}\right)$$

where subscripted variables and functions correspond to values in period t and 0.

Our study decomposes measured productivity into distinct components that recognize the stochastic nature of agricultural production. The decomposition analysis is not intended to explain the underlying drivers of productivity growth.[2] Instead, its intent is to disentangle weather-related growth components from more traditionally recognized sources of growth such as technological improvement, input-scale adjustments, and technology adoption. The decomposition procedure used adapts measures developed in Chambers and Pieralli (2020) to an annual framework. Their procedure, in turn, adapts measures developed by Caves, Christensen, and Diewert (1982), Färe, Grosskopf, and Lovell (1994), Kumar and Russell (2002), and Henderson and Russell (2005) to accommodate the stochastic nature of agricultural production. It starts by recognizing that production in some regions may lag technologically behind other regions. The relatively slow geographic diffusion of agricultural technologies across producers is a well-known empirical phenomenon (Schultz 1947; Griliches 1960; Gardner 2002). If a region operates inside the "best attainable" frontier, that departure from the frontier can be measured by dividing observed output, y_t, by the maximal output obtainable from its observed aggregate input and weather at time t, $f_t(x_t, w_t)$. We, therefore, define *efficiency* at time t as $E_t(y_t, x_t, w_t) = y_t/f_t(x_t, w_t)$.[3] Using this definition, the productivity index between period t and period $t-1$ decomposes into two components as

---

[2] An extremely large literature already exists on examining the drivers of agricultural productivity growth. Roughly put, those analyses take measured productivity growth and relate it via regression and other statistical analyses to "causal" factors such as education and investment in research and development. The *locus classicus* is Alston, Norton, and Pardey (1995), to which we refer the interested reader for an explanation of methods and summary of results.

[3] Our efficiency measure can also be recognized as an output-oriented radial distance function.

$$(2)\quad TFP(t,t-1)=\frac{E_t(y_t,x_t,w_t)}{E_{t-1}(y_{t-1},x_{t-1},w_{t-1})}\cdot\frac{f_t(x_t,w_t)/x_t}{f_{t-1}(x_{t-1},w_{t-1})/x_{t-1}}.$$

The first component, $\frac{E_t(y_t,x_t,w_t)}{E_{t-1}(y_{t-1},x_{t-1},w_{t-1})}$, measures the efficiency with which the technology available in period t was used relative to that with which it was used in period $t-1$. If this *efficiency index* exceeds one, production has moved closer to the frontier between the two periods. If it is less than one, production has moved further away. In the former case, a better job of incorporating available technological innovations into practice is done in period t than in period $t-1$, and in the latter, a worse job.

The second component of expression (2), $\frac{f_t(x_t,w_t)/x_t}{f_{t-1}(x_{t-1},w_{t-1})/x_{t-1}}$, measures the average product of x_t for realized weather w_t for frontier production $f_t(x_t, w_t)$, relative to the same measure for period $t-1$. If production were efficient in both periods, this ratio would measure actual productivity. Otherwise, it measures productivity that would be realized if production occurred on the efficient frontier.

We distinguish between three subcomponents of $\frac{f_t(x_t,w_t)/x_t}{f_{t-1}(x_{t-1},w_{t-1})/x_{t-1}}$. They are, respectively, an index of *technological change*, a *weather index,* and a *scale/input index*. The resulting decomposition is written as:

$$(3)\quad \frac{f_t(x_t,w_t)/x_t}{f_{t-1}(x_{t-1},w_{t-1})/x_{t-1}}=T_{t,t-1}W_{t,t-1}X_{t,t-1}$$

where

$$(4)\quad T_{t,t-1}=\left[\frac{f_t(x_{t-1},w_{t-1})}{f_{t-1}(x_{t-1},w_{t-1})}\frac{f_t(x_t,w_t)}{f_{t-1}(x_t,w_t)}\right]^{1/2}$$

is the index of technological change.

$$(5)\quad W_{t,t-1}=\left[\frac{f_t(x_{t-1},w_t)}{f_t(x_{t-1},w_{t-1})}\frac{f_t(x_t,w_t)}{f_t(x_t,w_{t-1})}\frac{f_{t-1}(x_{t-1},w_t)}{f_{t-1}(x_{t-1},w_{t-1})}\frac{f_{t-1}(x_t,w_t)}{f_{t-1}(x_t,w_{t-1})}\right]^{1/4}$$

is the weather index, and

$$(6)\quad X_{t,t-1}=\left[\frac{f_t(x_t,w_t)}{f_t(x_{t-1},w_t)}\frac{f_t(x_t,w_{t-1})}{f_t(x_{t-1},w_{t-1})}\frac{f_{t-1}(x_t,w_t)}{f_{t-1}(x_{t-1},w_t)}\frac{f_{t-1}(x_t,w_{t-1})}{f_{t-1}(x_{t-1},w_{t-1})}\right]^{1/4}x_{t-1}/x_t$$

is the scale index. The reasoning behind these decompositions is developed in Chambers and Pieralli (2020).

The index of technological change is the geometric average of two distinct measures of technological change. One, $\frac{f_t(x_{t-1},w_{t-1})}{f_{t-1}(x_{t-1},w_{t-1})}$, measures the shift in the production function between period $t-1$ and period t while holding x and w equal to that observed in period $t-1$. The other, $\frac{f_t(x_t,w_t)}{f_{t-1}(x_t,w_t)}$, measures the shift in the production function between period $t-1$ and period t while holding x and w equal to that observed in period t. The weather index is a geometric average of four weather indexes: two weather indexes computed relative to the period t technology holding x equal either to that observed in period t or to that observed in period $t-1$ and two weather indexes computed relative to the period $t-1$ technology holding x equal either to that observed in period t or to that observed in period $t-1$.[4]

[4] The geometric averaging of different indexes addresses the fact that *all indexes* (just as partial derivatives) depend upon the base at which they are evaluated. The most familiar examples, of course, are the differences between the Laspeyres and the Paasche indexes, which are addressed by calculating the Fisher Ideal Index as the geometric average of the two. The fact that productivity accounting requires choosing a base implies that all decomposition exercises will be path dependent (Kumar and Russell 2002). As just one example, both $\frac{f_t(x_{t-1},w_{t-1})}{f_{t-1}(x_{t-1},w_{t-1})}$ and $\frac{f_t(x_t,w_t)}{f_{t-1}(x_t,w_t)}$ exactly measure technological change but at different points on the frontier. Generally, they will differ,

The scale index is a geometric average of four measures of the average productivity of the aggregate input computed with respect to different technologies and different observed weather patterns.

The use of geometric averaging in constructing these components follows an approach established by Fisher (1922). More recently, however, Aczél (1990) has shown that the geometric average is the only relative merging procedure satisfying multiplicativity, positive homogeneity, and symmetry. Diewert and Fox (2017) show that the geometric average is a positively homogeneous symmetric mean that satisfies time reversal.[5]

These indexes are then translated into change form by taking logarithms to approximate percentage changes between period t and period $t-1$. Thus, measured productivity change decomposes as follows:

$$(7)\quad \begin{aligned} \ln TFP(t,t-1) &= \Delta T_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) \\ &+ \Delta W_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) \\ &+ \Delta X_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) \\ &+ \Delta E_{t,t-1}(y_t,y_{t-1},w_t,w_{t-1},x_t,x_{t-1}), \end{aligned}$$

where the *technological-change indicator* is

$$(8)\quad \begin{aligned} &\Delta T_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) \\ &= \frac{1}{2}[\ln f_t(x_t,w_t) - \ln f_{t-1}(x_t,w_t) \\ &+ \ln f_t(x_{t-1},w_{t-1}) - \ln f_{t-1}(x_{t-1},w_{t-1})], \end{aligned}$$

the *weather-change indicator* is

$$(9)\quad \Delta W_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) = \frac{1}{4}\begin{bmatrix} \ln f_t(x_{t-1},w_t) - \ln f_t(x_{t-1},w_{t-1}) + \ln f_t(x_t,w_t) - \ln f_t(x_t,w_{t-1}) \\ + \ln f_{t-1}(x_{t-1},w_t) - \ln f_{t-1}(x_{t-1},w_{t-1}) + \ln f_{t-1}(x_t,w_t) \\ - \ln f_{t-1}(x_t,w_{t-1}) \end{bmatrix},$$

the *input-change indicator* is

$$(10)\quad \Delta X_{t,t-1}(w_t,w_{t-1},x_t,x_{t-1}) = \ln x_{t-1} - \ln x_t + \frac{1}{4}\begin{bmatrix} \ln f_t(x_t,w_t) - \ln f_t(x_{t-1},w_t) + \ln f_t(x_t,w_{t-1}) \\ - \ln f_t(x_{t-1},w_{t-1}) + \ln f_{t-1}(x_t,w_t) - \ln f_{t-1}(x_{t-1},w_t) \\ + \ln f_{t-1}(x_t,w_{t-1}) - \ln f_{t-1}(x_{t-1},w_{t-1}) \end{bmatrix}$$

and the *efficiency-change indicator* is

$$(11)\quad \begin{aligned} &\Delta E_{t,t-1}(y_t,y_{t-1},w_t,w_{t-1},x_t,x_{t-1}) \\ &= \ln E_t(y_t,x_t,w_t) - \ln E_{t-1}(y_{t-1},x_{t-1},w_{t-1}). \end{aligned}$$

The empirical approach that we use to approximate the frontier production technology is nonparametric productivity analysis or data envelopment analysis (DEA). As applied by numerous authors in a variety of applied contexts (for example, Charnes et al. 1985; Byrnes et al. 1988; Fawson and Shumway 1988; Färe, Grosskopf, and Lee 1990; Färe et al. 1993; Chambers and Lichtenberg 1996; Kumar and Russell 2002; Henderson and Russell 2005; Murty, Russell, and Levkoff 2012; Chambers, Serra, and Lansink 2014), DEA builds upon methods originally developed by Afriat (1972) and traceable through Farrell (1957) to Koopmans' (1951) fundamental activity-analysis model. The essential idea is to use observed data to develop an approximation to the "best attainable technology" by enveloping it and then imposing sufficient structure to ensure that it is consistent

and geometric averaging ensures both measures are incorporated into a single measure of technological change. Constraining the technology by assuming a particular functional form, for example Cobb–Douglas exhibiting Hicks-neutral technical change, can eliminate the dependence but only at the expense of imposing arbitrary functional restrictions.

[5] The problem of index path dependence is well-known. Kumar and Russell (2002) provide a nice synopsis while Chambers and Pieralli (2020) provide further details for the decomposition method used in this paper. We take an agnostic approach in that our averaging procedure considers all possible paths that are available for the cross-period comparison.

with non-increasing returns to scale. Formally, this is done by taking an observed set of data on observed inputs and outputs, constructing the convex hull of the observed data, and then imposing "disposability properties" on that convex hull that are associated with the traditional notions of non-negative marginal returns, positive marginal costs, and non-increasing returns.

The data (described below) that we use to construct the empirical approximation to the technology consist of observations on aggregate output, aggregate input, and two measured weather variates for thirty-two Australian production regions for 1978 to 2013. Letting (y_{tk}, x_{tk}, w_{tk}) denote the observed values of aggregate output, aggregate input, and the two-vector of observed weather variates for region k at period t, the DEA approximation to the aggregate production function at period t for input bundle (x, w) is

$$f_t(x,w) = \max\left\{ \begin{array}{c} \sum_{j=1}^{t}\sum_{k=1}^{32}\lambda_{jk}y_{jk} : w = \sum_{j=1}^{t}\sum_{k=1}^{32}\lambda_{jk}w_{jk}, x \geq \sum_{j=1}^{t}\sum_{k=1}^{32}\lambda_{jk}x_{jk}, \\ 1 \geq \sum_{j=1}^{t}\sum_{k=1}^{32}\lambda_{jk}, \lambda_{jk} \geq 0, j=1,\ldots,t \end{array} \right\}. \tag{12}$$

This DEA approximation to the production function for a given (x, w) input vector is calculated as a linear program that chooses mixture terms (λ_{jk}, $j = 1, \ldots, t$, $k = 1, \ldots, 32$) to ensure that the aggregate output associated with (x, w) lies on the frontier of the "best attainable" technology. The task of the mixture terms is to select the combination of observed input and output variates that form the empirical approximation to the empirical frontier in the neighborhood of (x, w). Mathematically, the task of the inequality sign associated with the x constraint is to impose free disposability of x while the equality associated with the w constraints imposes weak disposability. This specification also satisfies Diewert's (1980) "sequential production set formulation" that ensures that technical know-how for a given set of inputs (including weather, climate conditions, etc.) does not degrade. A more thorough discussion of the method used to construct the DEA approximation to $f_t(x, w)$ is available in Chambers and Pieralli (2020) and in online supplementary appendix A.

The Data

The data used are aggregate input, aggregate output, and weather variates measured for thirty-two regions (from three different production zones) in Australian broadacre agriculture between 1978 and 2013.[6] The region-level data for aggregate input and aggregate output, constructed using standard productivity accounting practices, were obtained from Australian official sources (ABARES 2018).[7] Our weather variates are a *soil-moisture index* and a *growing degree-day (GDD)* measure. Total precipitation and average temperature for growing seasons are often used to represent weather-related effects. But they do not capture the role that moisture and temperature play in affecting plant and animal growth for different agro-ecological systems because air pressure, evaporation, runoff, soil attributes, heat accumulation, and other factors determine moisture delivered to the plant. The soil-moisture index and the GDD measure are preferred on this basis. We emphasize, however, that these variates only measure naturally occurring (beyond direct producer control) inputs. Alone, they do not measure

[6] A reviewer raises the important question of whether it is legitimate to use a common technology as representative of all three production (climatic) zones. The standard practice in productivity accounting is to define the technology as "all possible input and output combinations that are technically feasible." Thus, in principle, the technology is actually the metafrontier of all potentially available production frontiers in the Salter (1960) sense so that all observations should share a common frontier. This is true whether comparisons are done at the individual, village, regional, state, national, or international levels. Thus, in the body of the text we only report results generated from the metafrontier approach. Nevertheless, as a robustness check, we also performed the analysis using approximations of the metafrontier that were generated for each of the three production zones. These results are discussed and summarized briefly in online supplementary appendix B. They support the robustness of the results reported in the body of the paper.

[7] Sheng, Yang, and Zhao (2018) and ABARES (2018) contain a detailed discussion of the procedures used in calculating these TFP measures, and the aggregate inputs and outputs. In particular, they account for the differences in land quality across regions and their changes over time by using a hedonic approach when they estimate the land input.

Table 1. Descriptive Statistics on TFP Index, Soil Moisture Index, and Growing Season Degree Days: 1979–93 versus 1999–2013

	TFP index		Soil moisture index		Growing-season degree days	
	Mean/std	Skew/Kurt test	Mean/std	Skew/Kurt test	Mean/std	Skew/Kurt test
1979–93	1.047	16.100	0.227	23.460	4392.585	33.230
	(0.380)	(0.000)	(0.089)	(0.000)	(1043.092)	(0.000)
1999–2013	1.377	21.410	0.227	27.830	4434.941	33.380
	(0.516)	(0.000)	(0.087)	(0.000)	(1031.328)	(0.000)

Note: The skewness/kurtosis test for normality has been conducted using the method proposed by D'Agostino, Belanger, and D'Agostino Jr. (1990) and Royston (1991). Numbers in the parentheses below "Mean/std" are standard deviations while numbers in parenthesis below "Skewness/Kurtosis" test are p values.

the frequency or severity of droughts. That requires an intertemporal analysis of the evolution of moisture and GDD variates over time.[8]

The soil-moisture index was obtained from the Australian Water Availability Project (AWAP). That project uses model-data fusion methods that combine measurements and model predictions (Raupach et al. 2005a; Raupach et al. 2005b; Raupach et al. 2006; Trudinger et al. 2007; Trudinger et al. 2008) into a single soil-moisture index. The index ranges from −4 to +4, with −4 indicating extremely dry and +4 indicating extremely wet. It accounts for both historical and present soil moisture and water fluxes. The index is spatially and transtemporally comparable so that it can be used across different climatic zones (humid or arid). It is available for the entire Australian continent at a spatial resolution of 5 square km.

The growing season degree-day measure was calculated following the procedures used in Schlenker, Hanemann, and Fisher (2005) and Deschênes and Greenstone (2007). A base and a ceiling for daily average temperature determine the temperature thresholds. Specifically, a day with a mean temperature below 8°C contributes 0 degree days; a day with mean temperature between 8 and 32°C contributes the difference between mean temperature and 8 degree days (for example, a day with a mean of 15 contributes 7 degree days); a day with mean temperature above 32 degrees contributes 24 degree days. The growing season degree-day measure is calculated by adding the daily measures over two cropping windows: one period from 1st April to 31st October for the "winter season" and one period from 1st November to 31st March for the "summer season."

Both weather variates were calculated at the farm level (based on the Australian Agricultural and Grazing Industry Survey sample) by matching the location of 2,023 weather stations (or the map for soil moisture) with each farm. A weighted average of these measures for all farms falling inside each region is then used to compute the region-level measure with weights determined by sowing areas. Regional level weather measures are first calculated on a monthly base for the thirty-two regions between 1978 and 2013, and then aggregated using the two cropping windows defined above. The cropping/pasture areas for each growing season are used as weights.

Table 1 presents descriptive statistics on the empirical distribution of both measured weather variates. Figures 3 and 4 depict the resulting smoothed kernel densities of the weather variates for the 1979–93 and 1999–2013 periods. The Millennium-Droughts period brought both hotter and drier conditions. The average number of GDD increased between the periods. As figure 3(a) illustrates, the soil-moisture distribution shifted slightly to the left and became more leptokurtic. This reflects a movement toward slightly drier average conditions that are more tightly concentrated around the mean. Figures 3(b) and 3(c) overlay the smoothed kernel densities for each subperiod over histograms color coded (or dot-line framed) to discriminate mass located in each of the three zones (High-Rainfall, Wheat-Sheep, and Pastoral). Both panels demonstrate that the bulk of the mass for low-moisture levels falls in the Pastoral zone, the bulk of the mass for average moisture in the Wheat-Sheep zone, and the bulk of the mass for above average moisture in the High-Rainfall zone.

Figure 4(a) shows that the degree-day distribution shifted to the right and became more platykurtic in the Millennium-Droughts period.

[8] We thank an anonymous reviewer for emphasizing this point.

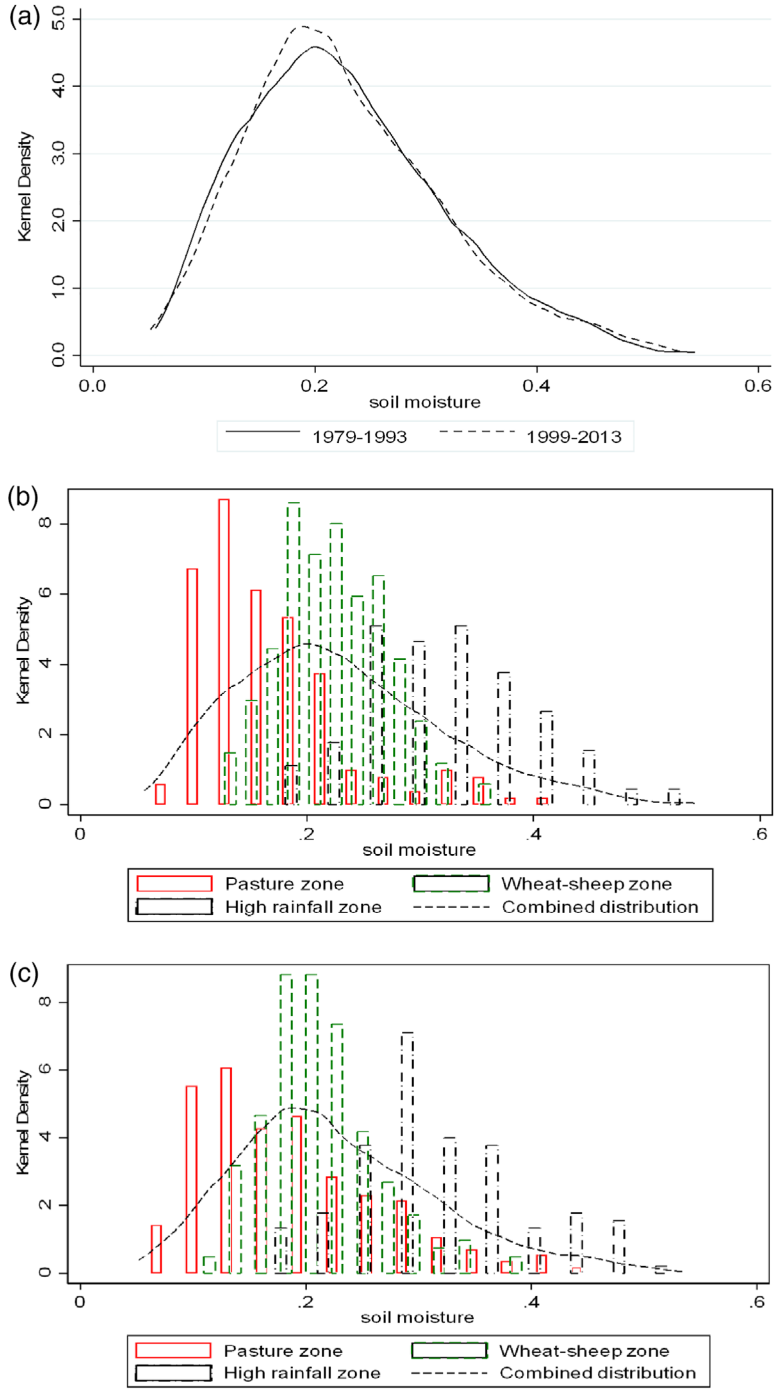

Figure 3. Smoothed kernel densities of Australian soil moisture in 1979–93 and 1999–2013 and histograms by climate zone. (a) Smoothed kernel densities of soil moisture index in 1979–93 and 1999–2013. (b) Smoothed kernel density of soil moisture and histograms by climate zone: 1979–93. (c) Smoothed kernel density of soil moisture and histograms by climate zone: 1999–2013

That pattern suggests a move toward warmer but more variable temperatures. Both the 1979–93 and 1999–2013 degree-day distributions give evidence of trimodality with mass concentrated near 4000 GDD, 5000 GDD, and 6500 GDD. As figures 4(b) and 4(c) demonstrate,

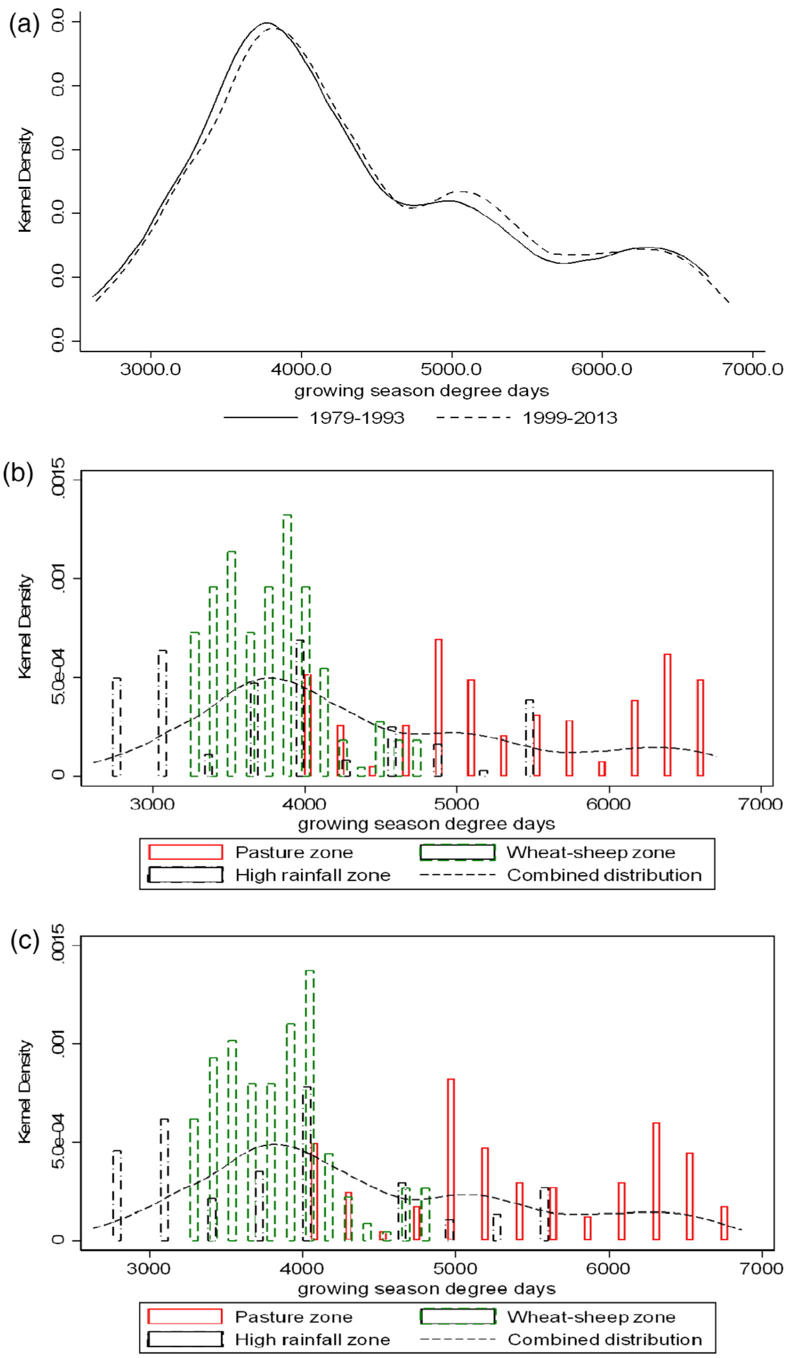

Figure 4. Smoothed kernel densities of Australian growing degree days in 1979–93 and 1999–2013 and histograms by climate zone. (a) Smoothed kernel densities of growing season degree days in 1979–93 and 1999–2013. (b) Smoothed kernel density of the growing season degree days and histograms by zone: 1979–93. (c) Smoothed kernel density of the growing season degree days and histograms by zone: 1999–2013

the trimodality closely parallels the different climatic zones. High-Rainfall areas exhibit the fewest growing season degree days, Wheat-Sheep areas intermediate levels, and the Pastoral zone accounts for the highest observed growing season degree days.

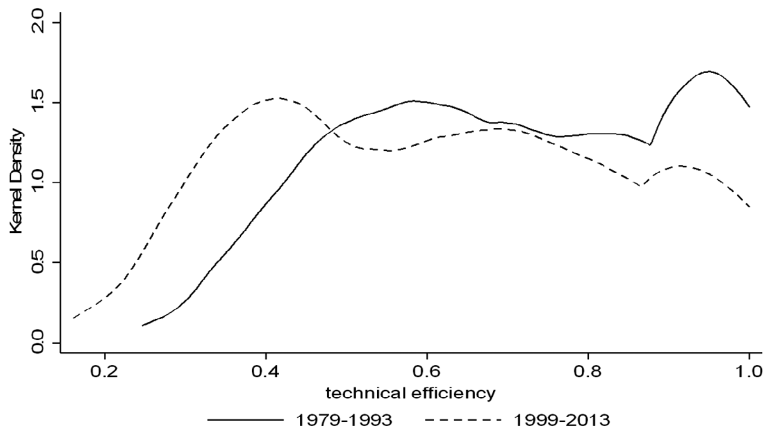

Figure 5. Smoothed kernel densities of technical efficiency between 1979–93 and 1999–2013

Empirical Results

Diffusion of Technology

Schultz (1947) argued that agricultural production could be both technically and allocatively inefficient. Griliches (1963) echoed these arguments in trying to explain observed patterns of US agricultural productivity growth. An oft-cited culprit was the relatively slow diffusion of agricultural technologies across and within regions (Griliches 1960; Gardner 2002; Mundlak 2005). For example, Griliches (1960) estimated that Iowan farmers had increased their hybrid-corn acreage anywhere from 10% to 90% more quickly than Alabaman farmers. We can examine the pattern of diffusion of agricultural technologies by seeing where different production units produce relative to the best attainable frontier. If diffusion of the technology is rapid and widespread, production units should be concentrated near the best attainable frontier. If diffusion is slow, significant concentrations of production units away from the frontier are to be expected.

In figure 5, we plot smoothed kernel densities for measured efficiency scores, $\frac{y_t}{f_t(x_t,w_t)}$, for the pre-Millennium-Droughts (PMD) period, 1979–93, and for the Millennium-Droughts (MD) period, 1999–2013. The relevant efficiency scores were derived from our empirical approximation to the "best attainable" technology. For the PMD period, the efficiency distribution appears bimodal with a mean of approximately 0.72. Mass is concentrated slightly below the mean and in the neighborhood of one. This suggests that productive regions fell into two groups during that time period. One, those in the neighborhood of one, operated relatively close to the frontier and maintained pace with the evolving technical frontier. Another lagged behind the moving frontier and benefited less from available technological developments.

The smoothed kernel densities for the MD period exhibit more mass concentrated in the lower end of the efficiency distribution and give evidence of trimodality. Average efficiency has clearly fallen. This suggests that with the appearance of the Millennium Droughts, diffusion and adoption of technological improvements both slowed and became more disparate across regions. Some mass is still concentrated in the neighborhood of one but clearly less than in the PMD period. What has emerged in its place is a group of intermediate laggards, with efficiency performance above the average but still falling away from the frontier. Instead of catching up with technological advances, this intermediate group appears to be falling behind.

Figure 6 provides another perspective on the nature of the change in the efficiency-score distribution between the two subperiods. The two panels there overlay the smoothed kernel densities with histograms color coded (or dot-line framed) according to climatic zones. In the PMD period, as illustrated by figure 6(a), a large percentage of the mass concentrated near one is associated with the High-Rainfall zone, followed by the Wheat-Sheep zone, and the Pastoral zone, respectively. The mid range of the distribution is dominated by High-Rainfall regions, while the lower tail of the distribution is dominated by Pastoral regions. Figure 6(b) shows that the trimodality observed in the MD period is associated with a

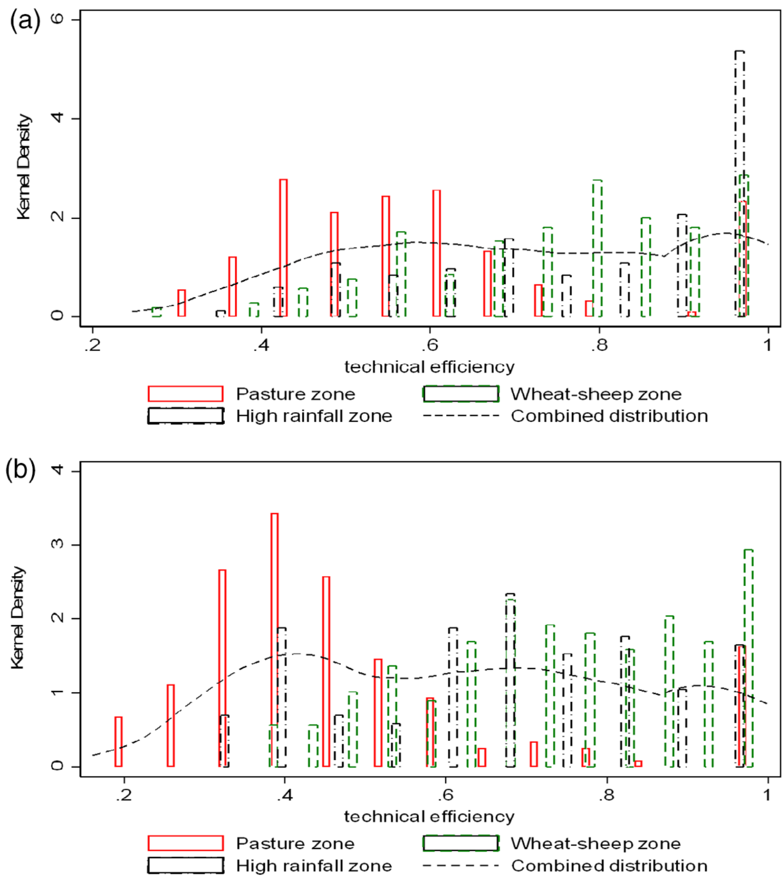

Figure 6. Smoothed kernel densities of technical efficiency scores in 1979–93 and 1999–2013, and histograms by climate zone. (a) Smoothed kernel density of efficiency scores and histograms by climate zone, 1979–93. (b) Smoothed kernel density of efficiency scores and histograms by climate zone, 1999–2013

number of High-Rainfall regions falling away from the "best attainable" frontier, while increasing numbers of regions in the Wheat-Sheep zone move into the neighborhood of one. The concentration of Pastoral regions in the neighborhood of one appears relatively unchanged, although Pastoral-zone scores are more concentrated in the lower half of the efficiency distribution. Moreover, the mass of High-Rainfall farmers falling in the lower tail of the efficiency distribution appears to have increased dramatically.

The story that emerges is that High-Rainfall zone farmers were less adaptable and did not maintain pace with the evolving best-practice technology during the MD period. There are varying explanations. One is that the input mix used in the High-Rainfall regions became less productive with lower moisture levels, while those for Wheat-Sheep regions improved, and Pastoral regions remain relatively unchanged. This, of course, is a story about adaptation to changing climatic conditions. It is well-documented, for example, that as US farmers moved westward and encountered changing climatic conditions, the process of adjusting to those changing conditions involved time-consuming experimentation with different varieties and input mixtures (Olmstead and Rhode 2011). In the Australian context, farmers encountered changing climatic conditions not as a result of geographic expansion, but because of changes in the severity and frequency of drought conditions.

The hypothesis that "adaptation to the frontier" depends upon climatic conditions warrants scrutiny. To that end, we examined the correlation between observed efficiency levels and the two weather variates for the two periods using bias-corrected regression techniques. These methods correct for the bounded [0, 1] nature of the support of the

Table 2. Bias-Corrected Regression: Efficiency Scores Related to Weather Variates

	1979–2013	1979–93	1999–2013
Dependent variable: Technical efficiency level			
Soil moisture index	0.579***	0.511***	0.674***
	(0.090)	(0.136)	(0.120)
Growing season degree days (log)	−0.521***	−0.476***	−0.518***
	(0.033)	(0.053)	(0.043)
Constant	4.857***	4.531***	4.782***
	(0.281)	(0.452)	(0.364)
Sigma	0.170***	0.158***	0.172***
	(0.005)	(0.007)	(0.007)
Number of observations	960	480	480
R-squared	0.238	0.209	0.184

Note: The bias-corrected regression refers to the routine in Kneip, Simar, and Wilson (2015). ***, ** and * denote significance at 1%, 5% and 10% level.

nonparametric efficiency scores (Kneip, Simar, and Wilson 2015). Three separate regressions were estimated: one for the entire sample period, 1979–2013, one for the PMD period, and one for the MD period. The results are reported in table 2. Estimated coefficients are uniformly significant at usual levels of confidence and broadly similar across the three specifications. These results should not be interpreted causally, but they do support the existence of a significant correlative relationship between the weather variates and adaptation to the frontier that seems stable across the different time periods. The results also suggest that measured efficiency is positively correlated with the soil-moisture variate and negatively correlated with the degree-day measure. Regions with higher soil-moisture and fewer degree days tend to operate closer to the "best attainable" frontier than regions with drier and hotter conditions. Similar results have been reported by others (Mendelsohn, Nordhaus, and Shaw 1994; Schlenker and Roberts 2009; Ortiz-Bobea, Knippenberg, and Chambers 2018; Chambers and Pieralli 2020).

Productivity Growth and its Components

We now examine *annual productivity-growth* performance in the PMD period and in the MD period. We calculated annual TFP change ($\ln TFP(t, t-1)$), technological change ($\Delta T_{t,\, t-1}$), weather change ($\Delta W_{t,\, t-1}$), input change ($\Delta X_{t,\, t-1}$), and efficiency change ($\Delta E_{t,\, t-1}$) for each of the thirty-two regions over the 1978–2013 period.

Figure 7 compares smoothed kernel density estimates of annual TFP change for 1979–93 and 1999–2013. Visually, the MD period exhibited a pattern of slower average productivity growth with more mass concentrated in its lower tail. This suggests that the upper support of the TFP distribution shifted to the left.

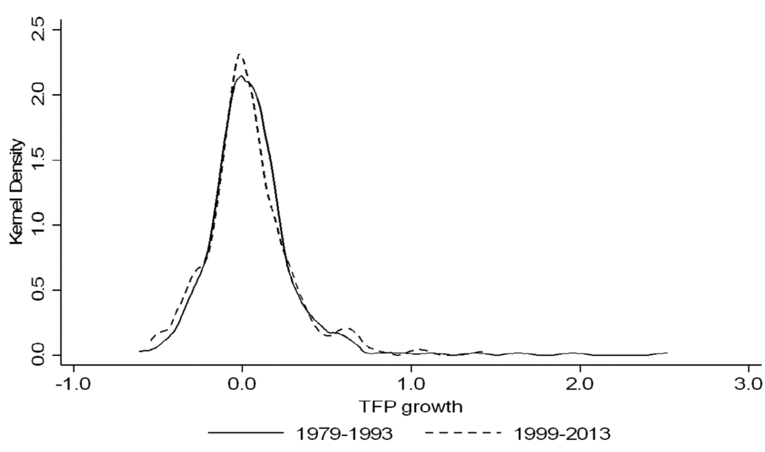

Figure 7. Smoothed kernel densities of annual TFP growth in 1979–93 and 1999–2013

Table 3. Kernel Analysis on TFP Growth and its Components between 1979–93 and 1999–2013

	TFP growth change	Technological progress	Input/scale adjustment	Efficiency change	Weather change
A (common shift)	−0.014**	0.204***	−0.001***	−0.007***	0.002
	(0.007)	(0.041)	(0.000)	(0.002)	(0.002)
D (dilation effects)	1.093***	5.660***	0.722	1.184***	0.541***
	(0.045)	(1.002)	(0.254)	(0.042)	(0.165)
S_left (left tail cut)	0.006	0.027	0.007	0.006	−0.082
	(0.007)	(0.069)	(0.031)	(0.005)	(0.051)
S_right (right tail cut)	0.001***	0.450***	−0.014	−0.003	0.037
	(0.000)	(0.039)	(0.030)	(0.004)	(0.048)
R-squared	0.792	0.998	0.967	0.871	0.816

Note: We conducted the kernel analysis following Combes et al. (2012). *, **, *** indicate statistical significance at the 10%, 5%, and 1% level. The total number of observations used in this exercise is 960 (480 × 2).

The MD distribution gives some evidence of bimodality. To augment the visual evidence, we have employed a nonparametric procedure due to Combes et al. (2012) that permits discrimination between three types of changes when comparing distributions: a mean or common shift, a dilation, and a support change.[9]

Table 3 reports the results from applying the Combes et al. (2012) procedure to the distributions of $\ln TFP(t, t-1)$, $\Delta T_{t,t-1}$, $\Delta W_{t,t-1}$, $\Delta X_{t,t-1}$, and $\Delta E_{t,t-1}$ for the PMD and MD subperiods. The first column refers specifically to $\ln TFP(t, t-1)$. The way to read each of the entries is that a positive (negative) estimate for the mean-shift parameter indicates a shift to the right (left) for the MD subperiod relative to the PMD subperiod, an estimate greater (less) than one for the dilation parameter indicates a more platykurtic (leptokurtic) distribution, a positive (negative) estimate for either a left or right support parameter means that the support has shifted toward the mean (away from mean). The reported R^2 indicates how closely the MD subperiod distribution fitted using the Combes et al. (2012) procedure approximates the observed distribution.

The statistical evidence reported in the first column of table 3 confirms the visual evidence on relative TFP performance. The estimated parameters for the mean shift, dilation, and upper support are all significant at the 0.05 level. The negative estimate for the mean shift indicates that average TFP growth was slower in the MD period than in the earlier period. The dilation parameter, which is greater than one, indicates greater variability in TFP performance as the Millennium Droughts hit. The positive parameter for the right support term indicates a leftward shift toward the mean. Thus, not only did average productivity growth slow, but some previously observed, high productivity growth patterns seem to have disappeared with the appearance of the Millennium Droughts.

Figure 8 depicts smoothed kernel density estimates for ΔT for the 1979–93 and 1999–2013 periods. Comparing the two distributions, the visual evidence suggests that a unimodal distribution of the PMD period was replaced by one containing evidence of bimodality. The latter bimodality appears to comprise a mass of regions concentrated very close to zero or little to no average technological improvement) and another (far smaller) mass of regions concentrated above the mean of the distribution. This pattern is evocative of what Quah (1996, 1997) has referred to as a "twin-peaks" phenomenon and suggests a divergence across regions in the rate of technological change. Some regions were continuing to innovate, while a larger group failed to innovate. The visual evidence also suggests that the upper support of the ΔT distribution appears to have shifted to the left suggesting a slowing rate of technological change for the most rapidly improving regions.

The appearance of the bimodal pattern in the ΔT distribution for the MD period deserves closer examination. In figure 9, we portray the smoothed kernel densities of ΔT overlaid with histograms color coded (or dot-line framed) according to climatic zones. The different panels show that the bimodal shape arises from Wheat-Sheep regions becoming more innovative. The leftward shift in the upper support of the ΔT distribution is associated with High-Rainfall regions exhibiting slower rates of ΔT (technological innovation).

[9] Technical details on the procedure are provided in an online supplementary appendix C.

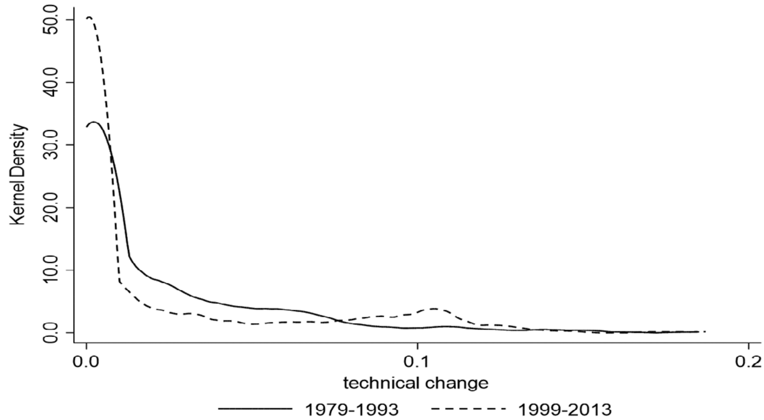

Figure 8. Smoothed kernel densities of technological change in 1979–93 and 1999–2013

Applying the Combes et al. (2012) decomposition analysis to the ΔT distributions (please see second column of table 3), we find significant statistical support for a rightward mean shift in the distribution, a large dilation of the distribution, and a leftward shift in the

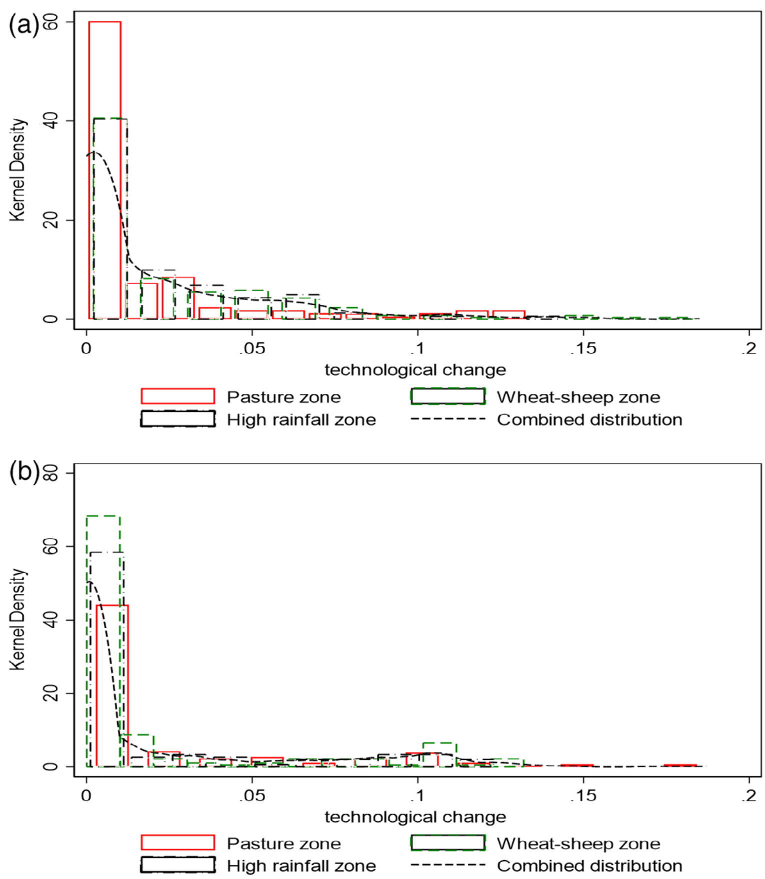

Figure 9. Smoothed kernel densities of technological change and histograms by zones in 1979–93 and 1999–2013. (a) Smoothed kernel density of technological change and histograms by zone, 1979–93. (b) Smoothed kernel density of technological change and histograms by zone, 1999–2013

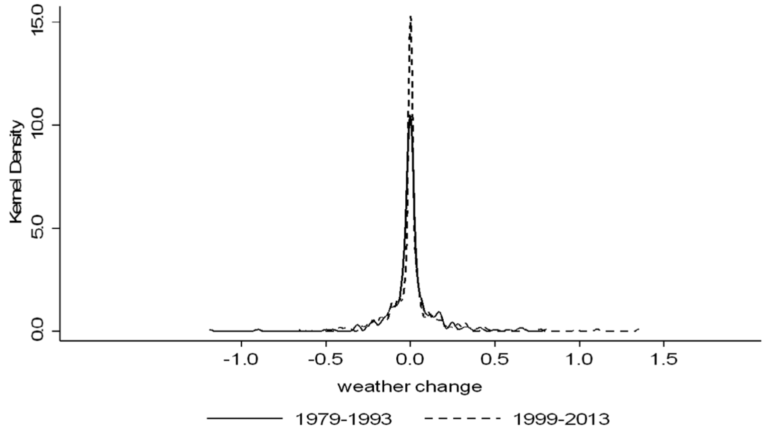

Figure 10. Smoothed kernel densities of weather change in 1979–93 and 1999–2013

upper support of the distribution. This suggests that growing interregional technological divergence accompanied the Millennium Droughts. High-Rainfall regions, which are dependent upon adequate moisture and moderate temperature, seemingly faltered in making technological innovations.

Both the ΔX and the ΔE distributions shifted leftwards with the onset of the Millennium Droughts. Both mean shift parameters reported in table 3 are negative and statistically significant at traditional confidence levels. As evidenced by the significant dilation factor reported in table 3, the ΔE distribution also became more platykurtic contributing to the observed increased variability in *TFP* growth between the two periods.

The final component of the decomposition is ΔW. Figure 10 presents its smoothed kernel densities for both periods. Both distributions are concentrated around a mean of zero. Thus, there is little evidence that the pattern of measured weather effects changed with the coming of the Millennium Droughts. On average, good years were roughly balanced by bad years in both periods so that mean effects were near zero. What is interesting, however, from the visual evidence in figure 10 is that the mass of observations concentrated in the neighborhood of zero has increased significantly between the two periods (the peak of the PMD distribution is lower than the peak of the MD one). The statistical evidence reported in table 3, with an estimated dilation parameter of 0.5, confirms the visual evidence of a more leptokurtic distribution.

Given the reported evidence (for example, Hughes, Lawson, and Valle 2017) that the Millennium Droughts caused large output losses, some may find a mean-zero ΔW effect in both periods counterintuitive. Intuition might suggest, instead, that the evidence would support a negative mean ΔW in the MD period. One should remember, however, how annually measured TFP responds to extreme weather events. When an extreme weather event occurs, annually measured TFP falls precipitously reflecting the sharp output drop relative to a stable input base. If weather returns to normal the following year, measured TFP rises dramatically as output "snaps back" to its normal level. The associated two-year average of ΔW should be about zero, even though a large output loss was encountered. In fact, if the input base were stable across the two periods the associated average measured rate of change would be positive. This happens because the fall in TFP percentage is measured against a higher output base than the following rise in TFP. Increasingly frequent drought conditions increase the frequency of these two-year, mean-zero events without changing the overall mean. They do, however, result in increased variability of measured TFP change.

The overall picture that emerges is slower and more variable *TFP* growth in the MD period. Importantly, the slowdown in average *TFP* growth is not associated with a slowing average rate of ΔT. In fact, the statistical analysis suggests that average ΔT rises in the MD period. The rightmost mode of the observed bimodal MD ΔT distribution contains sufficient mass to increase mean ΔT. But the lower average and more dispersed efficiency scores in the MD period suggest the diffusion of those technological advances slowed. At the same time that average ΔT increased, the upper tail of the empirical support of the ΔT shifted to the left. Fewer regions were making "extreme" technological improvements in the latter period. The evidence also suggests that the slowdown in *TFP* performance, the slowing technological diffusion, and the leftward shift upper support for the ΔT distribution are all associated with High-Rainfall regions falling away from the technological frontier. Regions in the Wheat-Sheep zones, however, seemed to have moved closer to it suggesting that their agricultural practices are most robust to extreme weather conditions.

A takeaway message, therefore, is that the sluggish *TFP* growth after 2000 does not result from technological failures. Instead, it seems more reflective of a slower and noisier process of diffusing technological advances. Two pieces of information are particularly noteworthy. The first is the confirmed link between extreme weather conditions and technology diffusion manifested by the bias-corrected regression analysis relating efficiency scores and the weather variates. The second is the visual evidence indicating cross-zone differences in how regions adapt to the frontier. During the PMD period, High-Rainfall regions were massed in the neighborhood of the technological frontier. They also often exhibited very high rates of productivity growth. Things changed in the MD period. Much of the failure to adapt to the evolving frontier was concentrated in the High-Rainfall zone. That zone has an agricultural infrastructure that was predicated on adequate moisture and moderate temperatures. Regions in such a zone necessarily will struggle with a short-term drought. This is expected and is one natural explanation for why measured agricultural *TFP* growth is periodically very negative. But when drought conditions become more the norm and not the exception, once-ideal input mixtures will need to adjust to more challenging environment. This is exactly what was required of American wheat farmers as they proceeded across the Western Plains (Olmstead and Rhode 2011). The evidence suggests that this process of adjustment and adaptation remains ongoing in the High-Rainfall zone.

Concluding Remarks

It is well known that yields and profitability are sensitive to extreme weather events (Schlenker, Hanemann, and Fisher 2005; Deschênes and Greenstone 2007; Schlenker and Roberts 2009; Tack, Barkley, and Nalley 2015). Less is known about how extreme weather impinges upon agriculture's sectoral performance. Evidence is scant, but what does exist suggests that climatic effects can be important (Liang et al. 2017; Ortiz-Bobea, Knippenberg, and Chambers 2018). We also know that developed-country annual agricultural productivity growth slumped dramatically around 2000 (Thirtle et al. 2004; Alston, Babcock, and Pardey 2010; Sheng, Mullen, and Zhao 2010; Alston, Andersen, and Pardey 2015; Sheng et al. 2020). The potential connection between the observed productivity growth slowdown and ongoing climate change begs to be examined, and Australian agriculture offers a natural laboratory.

We used a synthesis of different methods to investigate the interaction between climatic factors and Australian agriculture productivity. To incorporate and isolate weather's impact upon measured productivity performance, we modified Chambers and Pieralli's (2020) method for incorporating climate variates into a yearly productivity setting. Those results allowed us to decompose annual measured productivity growth into four separate components: technological change, input-scale adjustment, diffusion effects, and weather effects. We did this for thirty-two distinct production regions (in three different productions zones) for 1979–2013. These results were then partitioned into two subperiods, the PMD and the MD periods. We determined the subperiods using two criteria. First, we took the first period to end immediately prior to the first commonly acknowledged Millennium Drought. The starting point for the MD period was chosen to ensure that both periods had the same number of observations.[10] After the decompositions and partitioning were effected, we used nonparametric statistical analysis involving both estimation of smoothed kernel densities and kernel-analysis procedures due to Combes et al. (2012). These techniques can describe the behavior of various components of TFP change and can detect statistically discernible evidence of change in that behavior between

[10] Both the productivity comparisons and partitioning procedure used are distinct from those pursued in Chambers and Pieralli (2020). Their productivity change calculations were not annual but covered a thirty-year period, and their partitioning was chosen to treat the pre-PIK and post-PIK periods for the United States symmetrically.

the two subperiods. They cannot assess the causal nature of any such changes. That remains beyond the scope of our analysis.

The empirical results suggest that the slowdown in observed MD-period productivity growth is not statistically associated with a slowdown in average rate of technological change. Instead, weather-induced differences in patterns of technological diffusion, which are particularly concentrated in the High-Rainfall regions, seem to have played a more prominent role. In the PMD period, the High-Rainfall regions routinely developed and rapidly adopted technological improvements. But in the MD period, the pattern of their efficiency scores indicates that they struggled to incorporate technological advances into their production practices. They also exhibited lower rates of ΔT. This slowing of technological innovation and diffusion proved an important brake on their observed productivity performance. Whether it is a harbinger of a continued deterioration or reflective of needed further adjustments to changed climatic conditions is an open question needing further research. But it is of particular interest to note that Ortiz-Bobea, Knippenberg, and Chambers (2018) have recently reported similar results indicating that rain-fed US Midwestern agriculture's TFP performance is peculiarly sensitive to climate conditions.

Supplementary Material

Supplementary material are available at *American Journal of Agricultural Economics* online.

References

ABARES. 2018. Broadacre Agriculture. Available at: http://www.agriculture.gov.au/abares/research-topics/productivity/data (accessed 10 December 2017).

Aczél, Janos. 1990. Determining Merged Relative Scores. *Journal of Mathematical Analysis and Applications* 150(1): 20–40.

Afriat, Sydney N. 1972. Efficiency Estimation of Production Functions. *International Economic Review* 13: 568–98.

Alston, Julian M., Matthew A. Andersen, and Philip G. Pardey. 2015. The Rise and Fall of U.S. Farm Productivity Growth, 1910–2007. Staff Papers 200927, University of Minnesota, Department of Applied Economics.

Alston, Julian M., Bruce Babcock, and Philip G. Pardey. 2010. Shifting Patterns of Global Agricultural Productivity: Synthesis and Conclusion. In *The Shifting Patterns of Agricultural Production and Productivity Worldwide*, ed. Julian M. Alston, Bruce Babcock and Philip G. Pardey, 449–82. Ames, IA: The Midwest Agribusiness Trade Research and Information Center: Iowa State University.

Alston, Julian M., George W. Norton, and Philip G. Pardey. 1995. *Science under Scarcity: Principles and Practice for Agricultural Research Evaluation and Priority Setting*. Ithaca, NY: Cornell University Press.

Bureau of Meteorology (BOM). 2015. *Annual Australian Climate Statement 2015*. Available at: http://www.bom.gov.au/climate/current/annual/aus/2015/ (accessed 5 December 2019).

Byrnes, Patricia, Rolf Färe, Shawna Grosskopf, and C.A. Knox Lovell. 1988. The Effect of Unions on Productivity: U.S. Surface Mining of Coal. *Management Science* 3: 1037–53.

Cárdenas Rodríguez, Miguel, Ivan Haščič and Martin Souchier. 2016. "Environmentally Adjusted Multifactor Productivity: Methodology and Empirical results for OECD and G20 countries", OECD Green Growth Papers, No. 2016/04, OECD Publishing, Paris. https://doi.org/10.1787/5jlr2z7ntkf8-en.

Caves, Douglas W., Laurits R. Christensen, and W. Erwin Diewert. 1982. Multilateral Comparisons of Output, Input and Productivity Using Superlative Index Numbers. *Economic Journal* 92: 73–86.

Chambers, Robert G. and Erik Lichtenberg. 1996. A Nonparametric Approach to the von Liebig-Paris Technology. *American Journal of Agricultural Economics* 78(2): 373–86.

Chambers, Robert G., and Simone Pieralli. 2020. The Sources of Measured US Agricultural Productivity Growth: Weather, Technological Change, and Adaptation. *American Journal of Agricultural Economics* 1–29. https://doi.org./10.1002/ajae.12090.

Chambers, Robert G., Teresa Serra, and Alfons Oude Lansink. 2014. On the Pricing of Undesirable State-Contingent Outputs. *European Review of Agricultural Economics* 41(3): 485–509.

Charnes, Abraham, William W. Cooper, Boaz Golany, Lawrence M. Seiford, and Joel D. Stutz. 1985. Foundations of Data Envelopment Analysis for Pareto-Koopmans Efficient Empirical Production Functions. *Journal of Econometrics* 30(1/2): 91–107.

Combes, Pierre, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sebastien Roux. 2012. The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection. *Econometrica* 80(6): 2543–94.

D'Agostino, Ralph B., Albert Belanger, and Ralph B. D'Agostino, Jr. 1990. A Suggestion for Using Powerful and Informative Tests of Normality. *American Statistician* 44: 316–21.

Deschênes, Olivier, and Michael Greenstone. 2007. The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review* 97(1): 354–85.

Diewert, W. Erwin. 1980. Capital and the Theory of Productivity Measurement. *American Economic Review* 70: 260–7.

Diewert, W. Erwin, and Kevin J. Fox. 2017. Decomposing Productivity Indexes into Explanatory Factors. *European Journal of Operational Research* 256: 275–91.

Färe, Rolf, Shawna Grosskopf, and Hyunok Lee. 1990. A Nonparametric Approach to Expenditure-Constrained Profit Maximization. *American Journal of Agricultural Economics* 72: 574–81.

Färe, Rolf, Shawna Grosskopf, and C.A. Knox Lovell. 1994. *Production Frontiers*. Cambridge, Great Britain: Cambridge University Press.

Färe, Rolf, Shawna Grosskopf, C.A. Knox Lovell, and Suthathip Yaisawarng. 1993. Derivation of Shadow Prices for Undesirable Outputs: A Distance Function Approach. *Review of Economics and Statistics* 75(2): 374–80.

Farrell, Michael James. 1957. The Measurement of Productive Efficiency. *Journal of the Royal Statistical Society* Series A, 120: 253–90.

Fawson, Chris, and C. Richard Shumway. 1988. A Nonparametric Investigation of Agricultural Production Behavior for US Subregions. *American Journal of Agricultural Economics* 70: 311–7.

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. 2012. The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment. *American Economic Review* 102(7): 3749–60.

Fisher, Irving. 1922. *The Making of Index Numbers: A Study of their Varieties, Tests,and Reliability (No*. 1). New York: Houghton Mifflin.

Gardner, Bruce L. 2002. *American Agriculture in the Twentieth Century: How It Flourished, What It Cost*. Cambridge, MA: Harvard University Press.

Garnaut, Ross. 2011. *The Garnaut Review: Australia in the Global Response to Climate Change*. New York: Commonwealth of Australia and the Cambridge University Press.

Griliches, Zvi. 1960. Hybrid Corn and the Economics of Innovation. *Science* 132(3422): 275–80.

———. 1963. The Sources of Measured Productivity Growth: United States Agriculture 1940-60. *Journal of Political Economy* 71(4): 331–6.

Henderson, Daniel J., and R. Robert Russell. 2005. Human Capital and Convergence: A Production-frontier Approach. *International Economic Review* 46: 1167–205.

Hughes, Neal, Kenton Lawson, and Haydn Valle. 2017. Farm Performance and Climate: Climate-Adjusted Productivity for Broadacre Cropping Farms. ABARES Research Report 17.4 Department of Agriculture and Water Sciences, Canberra, Australia.

Kneip, Alois, Leopold Simar, and Paul W. Wilson. 2015. When the Bias Kills the Variance: Central Limit Theorems for DEA and FDH Efficiency Scores. *Econometric Theory* 31(2): 394–422.

Koopmans, Tjalling C. 1951. *Activity Analysis of Production and Allocation*. New York, NY: Wiley.

Kumar, Subodh, and R. Robert Russell. 2002. Technological Change, Technological Catch-up, and Capital Deepening: Relative Contributions to Growth and Convergence. *American Economic Review* 92: 527–48.

Liang, Xin-Zhong, You Wu, Robert G. Chambers, Daniel Schmoldt, Wei Gao, Chaoshun Liu, Yan-An Liu, Chao Sun, and Jennifer Kennedy. 2017. Determining Climate Effects on US Total Agricultural Productivity. *Proceedings of the National Academy of Sciences of the United States of America* 114(12): E2285–92.

Mendelsohn, Robert, William D., Nordhaus, and Daigee Shaw. 1994. The Impact of

Global Warming on Agriculture: A Ricardian Analysis. *American Economic Review* 84(4): 753–71.
Mundlak, Yair. 2005. Economic Growth: Lessons from Two Centuries of American Agriculture. *Journal of Economic Literature* 43(4): 989–1024.
Murty, Sushama, R. Robert Russell, and Steven B. Levkoff. 2012. On Modeling Pollution-generating Technologies. *Journal of Environmental Economics and Management* 64: 117–35.
Olmstead, Alan W., and Paul W. Rhode. 2011. Adapting North American Wheat Production to Climatic Challenges, 1839-2009. *Proceedings of the National Academy of Sciences of the United States of America* 108(2): 480–5.
Ortiz-Bobea, Ariel, Erwin Knippenberg, and Robert G. Chambers. 2018. Growing Climatic Sensitivity of U.S. Agriculture Linked to Technological Change and Regional Specialization. *Science Advances* 4(12): eaat4343.
Productivity Commission. 2009. Government Drought Support: Australian Government 2009. Productivity Commission, Report no. 46: Final Inquiry Report, Melbourne.
Quah, Danny T. 1996. Empirics for Economic Growth and Convergence. *European Economic Review* 40(6): 1353–75.
———. 1997. Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs. *Journal of Economic Growth* 2(1): 27–59.
Raupach, Michael R., Damian J. Barrett, Peter R. Briggs, and James Mac Kirby. 2005a. Simplicity, Complexity and Scale in Terrestrial Biosphere Modelling. In *Predictions in Ungauged Basins: International Perspectives on the State of the Art and Pathways Forward*, eds. Stewart W. Franks, Murugesu Sivapalan, Kuniyoshi Takeuchi and Yasuto Tachikawa, 239–74. Oxfordshire: IAHS.
Raupach, Michael R., Peter J. Rayner, Damian J. Barrett, Ruth S. DeFries, Martin Heimann, Dennis S. Ojima, Shaun Quegan, and Christiane C. Schmullius. 2005b. Model-Data Synthesis in Terrestrial Carbon Observation: Methods, Data Requirements and Data Uncertainty Specifications. *Global Change Biology* 11(3): 378–97.
Raupach, Michael R., Cathy M. Trudinger, Peter R. Briggs, and Edward A. King. 2006. *Australian Water Availability Project, CSIRO Marine and Atmospheric Research (Earth Observation) Component: Final Report*. Canberra: Australian Government Department of Agriculture, Fisheries and Forestry, CSIRO Marine and Atmospheric Research.
Royston, Patrick. 1991. Comment on sg3.4 and an Improved D'Agostino Test. *Stata Technical Bulletin* 1: 110–2. Reprinted in *Stata Technical Bulletin Reprints*, 3: 23–24. College Station, TX: Stata Press.
Salter, Wilfred E.G. 1960. *Productivity and Technical Change*. Cambridge: Cambridge University Press.
Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. 2005. Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach. *American Economic Review* 95(1): 395–406.
Schlenker, Wolfram, and Michael J. Roberts. 2009. Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change. *Proceedings of the National Academy of Sciences of the United States of America* 106: 15594–8.
Schultz, Theodore W. 1947. How Efficient Is American Agriculture? *Journal of Farm Economics* 29: 644–58.
Sheng, Yu, V. Eldon Ball, Carlos San Juan Mesonada and Kenneth Erickson. 2020. Testing the Embodiment Hypothesis and Productivity Convergence: Evidence from 17 OECD Countries between 1973 and 2011, CCAP Working Paper, No. 2020. 02,China Center for Agricultural Policy, Peking University.
Sheng, Yu, John D. Mullen, and Shiji Zhao. 2010. Has Growth in Productivity in Australian Broadacre Agriculture Slowed? Canberra: ABARES Technical Report. No. 2010.09.
Sheng, Yu, and Xinpeng Xu. 2019. The Productivity Impact of Climate Change: Evidence from Australia's Millennium Drought. *Economic Modelling* 76(1): 182–91.
Sheng, Yu, Sansi Yang and Shiji Zhao. 2018. Climate Change, Adaptation and Agricultural TFP: A Cross-regional Comparison of Australian Broadacre Farms. 62nd AARES Annual Conference Paper, Adelaide, Australia, 6–9 February 2018.
Stern, Nicholas. 2007. *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
Tack, Jesse, Andrew Barkley, and Lawton L. Nalley. 2015. Effect of Warming

Temperatures on US Wheat Yields. *Proceedings of the National Academy of Sciences of the United States of America* 112: 6931–6.

Thirtle, Colin, Lin Lin, Jim Holding, Lindie Jenkins, and Jenifer Piesse. 2004. Explaining the Decline in UK Agricultural Productivity Growth. *Journal of Agricultural Economics* 55(2): 343–66.

Trudinger, Cathy M., Michael R. Raupach, Peter J. Rayner, and Ian G. Enting. 2008. Using the Kalman Filter for Parameter Estimation in Biogeochemical Models. *Environmetrics* 19(8): 849–70.

Trudinger, Cathy M., Michael R. Raupach, Peter J. Rayner, Jens Kattge, Qing Liu, Bernard Pak, Markus Reichstein, Luigi Renzullo, Andrew D. Richardson, Stephen H. Roxburgh, Julie Styles, Ying Ping Wang, Peter R. Briggs, Damian J. Barrett, and Sonja Nikolova. 2007. OptIC Project: An Intercomparison of Optimization Techniques for Parameter Estimation in Terrestrial Biogeochemical Models. *Journal of Geophysical Research – Biogeosciences* 112(G2): G02027. https://doi.org/10.1029/2006JG000367.

Yang, Sansi, and C. Richard Shumway. 2016. Dynamic Adjustment in US Agriculture under Climate Change. *American Journal of Agricultural Economics* 98(3): 910–24.