Two Blades of Grass: The Impact of the Green Revolution

Douglas Gollin

Oxford University

Casper Worm Hansen

University of Copenhagen

Asger Mose Wingender

University of Copenhagen

We estimate the impact of the Green Revolution in the developing world by exploiting exogenous heterogeneity in the timing and extent of the benefits derived from high-yielding crop varieties (HYVs). We find that HYVs increased yields by 44% between 1965 and 2010, with further gains coming through reallocation of inputs. Higher yields increased income and reduced population growth. A 10-year delay of the Green Revolution would in 2010 have cost 17% of GDP (gross domestic product) per capita and added 223 million people to the developing-world population. The cumulative GDP loss over 45 years would have been US$83 trillion, corresponding to approximately one year of current global GDP.

> Whoever makes two ears of corn, or two blades of grass, to grow upon a spot of ground where only one grew before, would deserve better of mankind, and do more essential service to his country, than the whole race of politicians put together. (Jonathan Swift, *Gulliver's Travels*, part II, chapter VII)

This paper is a substantially revised version of earlier working papers with the same title. We thank the editor (Michael Greenstone), anonymous referees, Philipp Ager, Samuel

Electronically published June 3, 2021
[*Journal of Political Economy*, 2021, vol. 129, no. 8]

I. Introduction

How important is agricultural productivity growth in development? Early views of development assumed that most of the impetus for development and economic growth would necessarily come from the industrial sector, which was thought to offer the potential for rapid rates of productivity growth. In contrast, the agricultural sector in most developing countries was seen as backward and stagnant, with limited potential for growth (e.g., Rosenstein-Rodan 1943, Lewis 1951, and Nurkse 1953, echoed more recently in Matsuyama 1992). In recent years, agriculture's potential significance has been a theme in a renewed literature on structural transformation and economic growth. A new literature has offered theoretical models in which agricultural productivity growth is important for subsequent industrialization and in which agricultural productivity differences play a role in explaining cross-country disparities in income.[1] However, it has proved difficult to assess empirically the overall importance of agriculture's contributions to growth, and a lively policy debate remains on whether (and when, where, and how) governments should focus their development efforts on agriculture.

This paper contributes to the debate by studying how the Green Revolution affected economies in the developing world. The Green Revolution is arguably the most important episode of agricultural innovation in modern history and is best understood as an increase in agricultural productivity based on the application of modern crop-breeding techniques to the agricultural challenges of the developing world (Evenson and Gollin 2003a). New crop varieties were developed initially for rice, wheat, and maize; subsequently, scientists extended the Green Revolution technologies to a number of other crops. The increase in food production was

Bazzi, Derek Byerlee, Felipe Valencia Caicedo, Carl-Johan Dalgaard, Stefan Dercon, Dave Donaldson, Cheryl Doss, James Fenske, Nicola Gennaioli, Walker Hanlon, Peter Sandholt Jensen, Nicolai Kaarsen, Tim Kelley, Nathan Nunn, James Stevenson, Uwe Sunde, Tom Walker, and David N. Weil. We are also grateful for useful comments and suggestions from seminar participants at the University of Copenhagen, McGill University, the Paris School of Economics, the International Monetary Fund Research Department, the 7th Annual Workshop on Growth, History and Development at the University of Southern Denmark, the 2016 annual European Economic Association meeting in Geneva, the 2017 annual American Economic Association meeting in Chicago, the National Bureau of Economic Research conference on Trade and Agriculture, the Centre for Economic Policy Research conference Accounting for the Wealth of Nations: History and Theory, and the Brown University conference Deep Determinants of Economic Development. Early drafts of this paper were written while Gollin was serving on the Standing Panel on Impact Assessment (SPIA) of the CGIAR (formerly the Consultative Group on International Agricultural Research). Gollin's views do not represent those of SPIA or the CGIAR. Data are provided as supplementary material online.

[1] See, e.g., Gollin, Parente, and Rogerson (2002, 2007), Córdoba and Ripoll (2009), Restuccia, Yang, and Zhu (2008), and Vollrath (2011).

massive and nearly immediate in the irrigated rice-growing areas of Asia and the wheat-growing heartlands of Asia and Latin America. Other parts of the developing world received little benefit, however, from these early efforts—for reasons that are discussed in detail below.

How much did the Green Revolution matter? Did the advances in agricultural productivity generate large and long-lasting economic benefits? Answering this question poses obvious challenges for causal identification. Because growth in one sector of an economy will inevitably link to growth in other sectors, it is hard to find compelling evidence at the national level for the causal impacts of agricultural productivity growth. Using variation in productivity within countries, at a narrower geographic scale, several papers have made use of quasi-natural experiments (e.g., Hornbeck and Keskin 2014 and Bustos, Caprettini, and Ponticelli 2016) or structural estimation (e.g., Foster and Rosenzweig 2004, 2007) to look at the cross-sectoral impacts of changes in agricultural productivity. However, these local effects can be difficult to extrapolate to full general equilibrium impacts on aggregate economies. In poor countries with large fractions of their workers in agriculture, the main mechanisms of structural transformation are not played out within local labor markets. Instead, they often involve large-scale movements of people across locations—from rural areas to cities or from one region to another. Studies that emphasize the local movements of people will miss these broader and more secular changes.

Informed by a theoretical model, we estimate the impact of the Green Revolution on national economies in two steps. First, we leverage variation in the global diffusion of HYVs (high-yield varieties) in a staggered-adoption design to estimate the impact on crop yields. For each crop, we are able to use historical records on the breeding and release of HYVs to identify a specific release date at which the new Green Revolution technology became available to the developing world. Given these release dates, we compare yields of crops for which HYVs became available to yields of crops that did not benefit (or had not yet benefited) from comparable varietal improvement research. In a sample of 90 developing countries, we find that HYVs increased annual yield growth of some crops by as much as 1.3 percentage points, and we demonstrate that this difference-in-difference result is not driven by preexisting trends in crop yields. Wheat and rice experienced the highest yield increases; other important crops, such as cassava and sorghum, were less affected by the Green Revolution—both because HYVs became available at a later date and because HYVs had only a modest impact on yields. We use this variation in the second step of our analysis, in which we estimate how the Green Revolution affected economic growth, demography, and development more broadly. By combining our crop-specific estimates of the impact of HYVs with country-specific shares of each crop in total agricultural production before the Green Revolution, we construct a measure of the exogenous impact of HYVs on aggregate yields (for fixed allocations of land and labor). The resulting variable is similar to a shift-share (or Bartik) instrument. But contrary to most applications of shift-share research designs, which rely on observed aggregate trends to draw inference at the disaggregated level, our design uses the exogenous yield trends that we estimated in the first part of our analysis. This allows causal inference not just at the country level but also at the developing-world level, making it possible to quantify the global effects of the Green Revolution.

Our shift-share variable indicates that HYVs increased yields of food crops by 44% between 1965 and 2010. The total effect on yields is even higher because of substitution toward crops for which HYVs were available and because of reallocation of land and labor. Beyond agriculture, our baseline estimates show strong, positive, and robust impacts of the Green Revolution on different measures of economic development. Most striking is the impact on GDP (gross domestic product) per capita. Our estimates imply that delaying the Green Revolution for 10 years would have reduced GDP per capita in 2010 by US$1,273 (adjusted for PPP [purchasing power parity]), or 17%, across our full sample of countries. The dollar amount is large, in part because some of the countries grew relatively rich during the period we study: the comparable loss in today's least developed countries is US$392. By 2010, the cumulative global loss of GDP of delaying the Green Revolution 10 years would have been about US$83 trillion—roughly a year of present-day global GDP. Needless to say, this surpasses the amount of resources that went into developing HYVs by several orders of magnitude. The income loss would have been much greater had the Green Revolution never happened, perhaps reducing GDP per capita in the developing world to 50% of its current level, if our estimates are taken at face value—although we stress that this number is subject to considerable uncertainty and depends on a somewhat implausible counterfactual. Despite these reservations, the results of this paper clearly place the Green Revolution among the most important economic events in the twentieth century.

We find no evidence that the gains from increased agricultural productivity were offset by any Malthusian effects; the increased availability of food does not appear to have been eroded by population increases. Instead, we find a negative effect of the Green Revolution on fertility. Our estimates suggest that the world would have contained more than 200 million additional people in 2010 if the onset of the Green Revolution had been delayed for 10 years. Lower population growth increased the relative size of the working-age population, leading to a demographic dividend that accounts for roughly one-fifth of our estimated effect on GDP per capita. Our paper also sheds light on a concern, often expressed in the literature, that agricultural productivity improvements would pull additional land into agriculture at the expense of forests and other environmentally valuable land uses. We find evidence to the contrary: in keeping with the "Borlaug hypothesis," the Green Revolution tended to reduce the amount of land devoted to agriculture.[2]

A large literature considers the social, economic, and environmental impacts of the Green Revolution; it would be too ambitious to review this literature here. Recent surveys in the economics literature include Renkow and Byerlee (2010) and Pingali (2012).[3] Our paper addresses some of the same macro-scale questions that have previously been considered using models of varying structures and with differing assumptions; see, for example, Evenson and Rosegrant (2003) and Perez and Rosegrant (2015). A recent survey of these models can be found in Godfray and Robinson (2015). In contrast to these approaches, our analysis is based on econometric evidence and is in some respects closer to papers that combine spatial variation in geography with the arrival of new technologies whose impacts depend on geography, such as the potato (Nunn and Qian 2011), GM (genetically modified) crops (Bustos, Caprettini, and Ponticelli 2016), and fracking (Bartik et al. 2019). Our paper is perhaps closest to a small set of recent papers that similarly combine spatial variation with time variation to study the Green Revolution—specifically, works by Moscona (2019), Bharadwaj et al. (2020), and von der Goltz et al. (2020).

II. Origins of the Green Revolution

Although formal programs of scientific research on crop improvement in developing countries can be traced back into the nineteenth century, the timing of the initial Green Revolution and its subsequent patterns of diffusion were largely exogenous to individual countries. The argument we make is based on three claims. The first is that the initial Green Revolution technology was almost entirely developed in a set of international institutions that revolutionized crop breeding through large-scale crossing based on a then-modern understanding of genetics; these institutions also had access to a wide range of genetic material, having assembled large collections of traditional crop varieties that had not previously

[2] Norman Borlaug (1914–2009) was a wheat scientist closely associated with the early years of the Green Revolution. Borlaug won the Nobel Peace Prize in 1970 for his work in developing and promoting the Green Revolution, most notably through his efforts in wheat breeding. Borlaug argued forcefully that improved varieties and higher agricultural productivity would lead to reduced pressure on land resources, as higher production would be achieved through intensification rather than extensive expansion of agricultural area. This argument was dubbed the "Borlaug hypothesis" by Angelsen and Kaimowitz (2001, 3).

[3] This paper is also related to an even larger literature that considers the impact of agricultural science on economic and social outcomes at a more geographically limited scale. This literature is surveyed by Maredia and Byerlee (2000); more recent contributions include Thirtle, Lin, and Piesse (2003), Pingali and Kelley (2007), Dalrymple (2008), Raitzer and Kelley (2008), Rusike et al. (2010), and Costinot and Donaldson (2011).

been available to breeders. Our second claim is that the timing of the initial research was driven by a mixture of humanitarian and geopolitical concerns. In this sense, it was not driven by an assessment of the subsequent growth prospects of any particular country or set of countries. (If anything, the focus was on countries that seemed at risk of famine and political crisis.) The third is that the HYVs produced through Green Revolution research were made widely available in countries producing those crops. Because these technologies were developed in public sector institutions and made available in the public domain, and because HYV seeds were essentially self-replicating, the diffusion of the technology was not significantly limited or mediated by proprietary control or even by the capabilities of governments. Many of the Green Revolution HYVs diffused through farmer-to-farmer sales or sharing of seeds. This also meant that research targeted particular agronomic and phenotypic problems thought to have widespread relevance, rather than focusing on specific countries or on the most profitable market segments. Together, these three claims support the proposition that the differential impact of agricultural research on developing economies reflected factors substantially exogenous to those countries.

A. The Institutional Basis for the Green Revolution

Although many developing countries had some indigenous and colonial programs of crop improvement, it is a reasonable generalization to say that few developing countries had large or systematic programs of crop improvement before 1950. Colonial programs of agricultural research tended to focus on nonfood crops, such as sugar, that provided raw materials for industry or were consumed in the colonial heartland. Food crops tended to receive a low priority. To the extent that there were active programs of research on food crops, as in India, in the first half of the twentieth century, they tended to focus on identifying vigorous strains of existing varieties rather than developing new lines. Early Green Revolution technologies were closely linked to an institutional innovation in agricultural research that created a new set of plant-breeding institutions. In particular, the HYVs were closely associated with the creation of new internationally funded research centers (IARCs) and the large-scale mobilization of scientific resources. Because of this, the origins of the Green Revolution can be dated fairly precisely.[4]

[4] It is true that plant breeding took place before the Green Revolution within many national programs. But before the Green Revolution, varietal improvement in most national programs was heavily based on selection from existing varieties, rather than through "crossing" (or hybridization). To the extent that crossing took place, it was carried out on a small scale. For instance, the entire Indian national program in rice research appears to have been making no more than a few dozen crosses per year around 1960; by contrast, in the early

The earliest large-scale programs of international research took place in rice, wheat, and maize—the world's most important food crops. Following some early exploratory work in the 1940s and 1950s, the first of the Green Revolution institutions was created in 1960, in the form of the International Rice Research Institute (IRRI), located near Los Baños in the Philippines. In 1967, a sister institution was born: the International Center for Maize and Wheat Improvement (CIMMYT), with headquarters near Texcoco, Mexico. These two research centers were funded by a group of aid donors, including the Ford and Rockefeller Foundations, as well as a number of national aid agencies. CIMMYT grew out of an ongoing program of wheat research that the Rockefeller Foundation had been funding in Mexico since the late 1940s.[5] The history of the early Green Revolution has been documented in a number of sources, for example, Dalrymple (1974, 1985, 1986) and Barker, Herdt, and Rose (1985). Breeding efforts at these institutions were subsequently extended to other crops and other research centers, as discussed below.

In both rice and wheat, these early efforts reflected an emerging view that rich countries had both obligations and opportunities to encourage development in the newly independent countries of Africa and Asia, in the wake of the Second World War. This view coincided with geostrategic concerns triggered by the Cold War. The threat of agrarian revolutions in Asia and Latin America seemed to call for efforts to promote rural development (for a detailed discussion, see Perkins 1997). It was presumably not a coincidence that the United States, being pulled steadily into a war in Indochina and fearing a domino effect, chose to support investments in rice research or that it would support a wheat research program that was based in Mexico.

Against this backdrop, rice breeding began at IRRI in late 1962. Within the first weeks of breeding effort, scientists made a cross (designated IR8) that gave rise, after several years of further selection, to what would eventually prove to be the first "megavariety" of rice. The other initial research center, CIMMYT, began distributing HYVs of wheat and maize even before it was formally founded. In the wake of the successes of IRRI and CIMMYT, two additional centers were created in 1967: the International Institute for Tropical Agriculture (IITA) in Ibadan, Nigeria, and the International Center for Tropical Agriculture (CIAT) in Cali, Colombia. These institutions were assigned mandates for additional crops and

---

years of the International Rice Research Institute, breeders averaged over 2,500 crosses per year.

[5] IRRI, too, had a modest precursor program, a small breeding effort initiated under the auspices of the United Nations (UN) Food and Agriculture Organization (FAO) in the 1940s and 1950s. It is safe to say, however, that there was no large-scale or systematic effort to breed new rice varieties for the developing world before 1960.

different agroecologies; the subsequent rolling out of additional centers provides a valuable tool for identification in our analysis.[6]

To a degree, adaptive breeding—the effort to tailor HYVs to specific agroecological niches and to address problems of local importance—has been carried out by national governments through agricultural research systems, university-based research programs, and other local research. A concern for our identification strategy is that this effort may thereby reflect institutional capacity, raising the possibility that the diffusion curves for different countries are related to general institutional factors that might lead to growth through other channels. But what is clear is that even for the most advanced developing countries, adaptive breeding has continued, even to the present day, to rely heavily on research emerging from the CGIAR. Many or most HYV crops in the developing world continue to use genetic material that can be traced to the CGIAR, as documented by Evenson and Gollin (2003b) for the period through the 1990s. More-recent studies describe the continuing importance of international research for the diffusion of HYVs in sub-Saharan Africa (Walker and Alwang 2015) and South Asia (Pandey, Velasco, and Yamano 2015).

B. Timing of the Green Revolution across Crops

The start of the Green Revolution can be dated quite precisely. As noted above, the first high-yielding rice varieties were crossed in 1962 at IRRI, and after several generations of selection, they were initially released in 1965 to national research programs in rice-growing countries around the world. For wheat, it is similarly possible to identify a zero date for the Green Revolution: the first successful crosses from the Rockefeller wheat program took place in the 1950s, but they were not released to farmers in other developing countries until 1965. Maize followed soon after. For each crop, we can identify with reasonable precision the date at which the research institution first released a variety based on breeding work that took place within the institution. Table 1, compiled on the basis of historical records of varietal releases and other analysis of breeding data, shows the release dates of HYVs for different crops; detailed documentation is in table A1; tables A1–A13 are available online.

The public nature of the international agricultural research centers means that the HYVs they helped produce were made freely available, so the dates identified in table 1 reflect the year in which any developing

[6] There are now 15 such institutions that carry out agricultural research on subjects ranging from aquaculture to livestock science to water management. These centers operate collectively as an entity known as CGIAR (formerly known as the Consultative Group on International Agricultural Research), with an annual budget approaching $1 billion. Its research is funded by national and multilateral development agencies, nongovernmental organizations, private philanthropies, and other donors.

TABLE 1
Release Year of First HYV by Crop

Crop	Crop Type	IARC Research Center (Host Country)	Center Founding	Initial Crop Mandate	First Released Material	Year Used in Analysis
Barley	Cereals	ICARDA (Syria)	1977	1977	1979	1979
Maize	Cereals	CIMMYT (Mexico)	1963	1963	1966	1966
Millet	Cereals	ICRISAT (India)	1972	1972	1982	1982
Rice	Cereals	IRRI (Philippines)	1960	1960	1966	1966
Sorghum	Cereals	ICRISAT (India)	1972	1972	1983	1983
Wheat	Cereals	CIMMYT (Mexico)	1963	1963	1965	1965
Plantain	Fruit	IITA (Nigeria)	1967	1972	2002	. . .
Soybean	Pulses	IITA (Nigeria)	1967	Unclear	1979	. . .
Yam	Roots and tubers	IITA (Nigeria)	1967	1970	1990	1990
Dry beans	Pulses	CIAT (Colombia)	1970	1973	1979	1979
Cassava	Roots and tubers	CIAT (Colombia)	1970	1973	1984	1984
Potato	Roots and tubers	CIP (Peru)	1971	1971	1990	. . .
Sweet Potato	Roots and tubers	CIP (Peru)	1971	1988	Mid-1980s	1999
Groundnut	Pulses	ICRISAT (India)	1972	1974	1985	1985
Pigeon pea	Pulses	IITA (Nigeria)	1967	1972	2002	2002
Chickpea	Pulses	ICARDA (Syria)	1977	1977	1984	1984
Faba bean	Pulses	ICARDA (Syria)	1977	1977	1986	1986
Lentils	Pulses	ICARDA (Syria)	1977	1977	1980	1980
Cowpea	Pulses	IITA (Nigeria)	1967	1970	1974	1974

Note.—By crop, the year in which the first HYV was released, along with the IARC from which it originated. Our empirical analysis does not include potatoes, soybeans, and plantain, for reasons explained in n. 15. CIP = Centro Internacional de la Papa (International Potato Center); ICARDA = International Center for Agriculture Research in the Dry Areas; ICRISAT = International Crops Research Institute for the Semi-Arid Tropics. For further details, see table A1.

country could potentially have adopted HYVs of a given crop. Wheat is an exception to this rule and precisely for this reason provides an instructive example. Wheat breeding in Mexico initially focused on disease resistance—in particular, resistance to wheat leaf rust, a pathogen that significantly reduces yields. The first rust-resistant variety was released in 1948, and further improved varieties followed in the 1950s. Work on semidwarfism began in the mid 1950s and culminated with the release of the first truly high-yielding varieties in Mexico.[7] On the basis of the success of these varieties, Norman Borlaug, the lead scientist at the Mexican research center, brought seeds of the initial semidwarf wheats to India in 1963. In 1965,

[7] Semidwarf varieties of rice and wheat were central to the Green Revolution. Because they were shorter and had stiffer straws, they converted greater fractions of plant energy to grain and less to producing stalks and leaves. They also provided sturdier structural support for heavy grain production, which meant, in turn, that they produced well under high doses of fertilizer—whereas traditional (taller) varieties of rice and wheat had a tendency to topple (or "lodge") when fertilized intensively.

two semidwarf varieties originating from the Mexican program were released in India; more or less concurrently, semidwarf varieties were released in Pakistan. By 1970, nearly 10 million ha of HYV wheat had been planted in Bangladesh, India, Nepal, and Pakistan; by 1977–78, the area planted with HYVs had doubled and accounted for approximately two-thirds of the wheat area in those countries (Dalrymple 1985). This differential timing of the arrival of HYVs of wheat in Mexico and India is clearly visible in figure 1, which shows that wheat yields start to increase in Mexico in the 1950s and in India and other developing countries in the late 1960s.

The Mexican case is unique in the sense that the first HYVs were developed in a research program that did not yet have standing as an international institution. As a result, the diffusion of the wheat semidwarf varieties took place within Mexico slightly before the varieties became available in other countries. For all our other crops, HYVs developed at the international research centers became available to all countries at effectively the same moment—either upon a formal initial release from the international center or through the inclusion of the material in "nurseries" of promising experimental material that were shared with researchers across the developing world. The timing and magnitude of the Green Revolution is not the same in all countries, however, because HYVs of

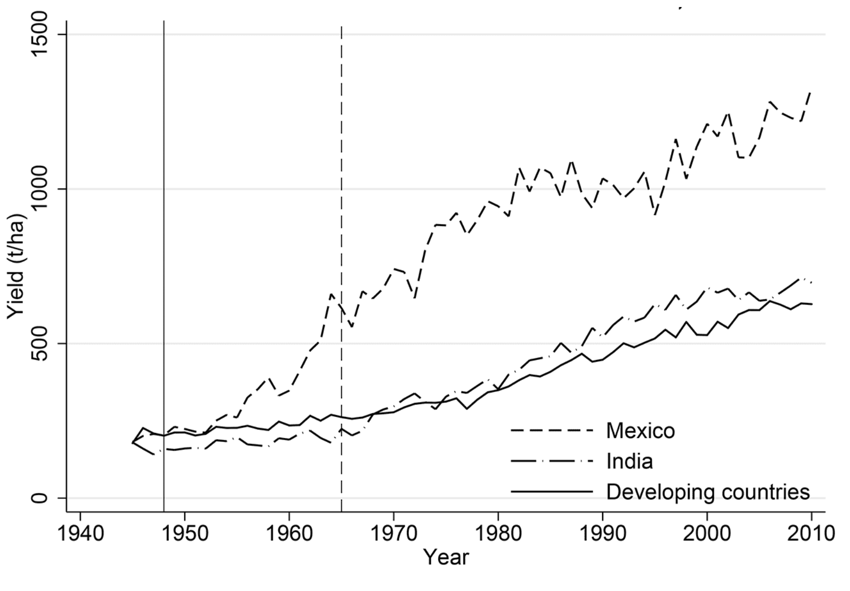

FIG. 1.—Wheat yields in Mexico, India, and the average country in the developing world (i.e., our baseline sample). The solid vertical line indicates the release date of the first HYV in Mexico. HYVs did not become available in other countries until the agricultural year 1965/66, when wheat HYVs were released in India and a number of other countries in Asia (dashed vertical line).

different crops became available at different times. The earliest releases of HYVs were in the three most important cereal crops (rice, wheat, and maize), whereas other cereal crops (barley, sorghum, and millet) saw little in the way of HYV development until the 1980s. High-yielding varieties of cowpeas (also known as black-eyed peas) were first released in the mid-1970s, but most other beans and legumes (lentils, chickpeas, etc.) did not see successful HYVs until the early 1980s. Successful HYVs for most root crops did not arrive until even later.

There are a number of plausible interpretations for the variation in timing of the Green Revolution across crops. To a large extent, the early successes in rice, wheat, and maize reflected the fact that advanced research institutions in developed countries had been working on these crops for decades before the beginning of the Green Revolution. Breeders could begin with elite lines from North America, Europe, and Japan. Moreover, they had a good understanding of the extent of genetic diversity and the sources of useful genes. The situation was different for tropical root crops (e.g., cassava and sweet potato) and for minor crops in rich countries, such as millet and sorghum. The development of HYVs of these crops required far more prior research. These differing initial stocks of knowledge and improved genetic materials create another source of exogenous variation in the timing and extent of the Green Revolution. In practice, this meant that countries that were heavily dependent on rice or wheat agriculture received an earlier—and potentially stronger—boost from the Green Revolution than those that relied on root crops or on other cereal grains, such as barley, sorghum, and millet.

III. Motivating Theory

To motivate our empirical analysis, this section provides a simple theoretical framework that connects HYV adoption at the farm level to reallocation within the agricultural sector and subsequently to economy-wide outcomes.[8] Consider a country with N regions of fixed size $\bar{X}_k$, $k = 1, 2, \ldots, N$. Regions are not necessarily geographically distinct, but they may refer to different ecologies in the same area (e.g., hills and valleys). At time t, a share of land within each region, $s_{kt} \in [0, 1]$, is used for crop production; the rest is left fallow. Let $X_{kt} = s_{kt}\bar{X}_k$ denote cropped area in region k. To simplify, assume that regions have distinct agroecologies such that region k can grow only a single crop k. Land within each region is divided into infinitesimally small plots, indexed by $i \in [0, \bar{X}_k]$, that are heterogeneous in terms of soil quality. Let a_{ikt} denote soil productivity of plot i in region k at time t, and let plots be indexed according to their

[8] We report only the key equations, but the intermediate steps are provided in app. B; apps. A–I are available online.

productivity levels such that a_{ikt} is decreasing in i. The productivity parameter a_{ikt} is time varying because it depends on the available technology, climate change, and so on. We assume that the most productive soils are planted first, in a Ricardian sense, meaning that crops are grown in region k on all plots $i \leq X_{kt}$ and that all plots $i > X_{kt}$ are left fallow. Moreover, let a_{ik} be distributed smoothly and differentiably across plots in such a way that the average productivity of cropland within a region is given by

$$\frac{1}{X_{kt}} \int_0^{X_{kt}} a_{ikt}\, di = a_{kt} s_{kt}^{-\delta}, \quad 0 < \delta < 1, \tag{1}$$

where a_{kt} is a region-specific (and by implication crop-specific) productivity parameter to be specified below. The term $s_{kt}^{-\delta}$ captures that average soil productivity declines with the share of land devoted to the single crop, irrespective of region size.

Let the unit of production in agriculture be a family farm. Each family owns one plot of land and supplies one unit of labor, which is the only other input in production besides land. The mass of farms in region k is consequently equal to the quantity of farm labor in region k, which we denote L_{kt}. These assumptions, along with equation (1), imply that aggregate agricultural output in region k is a Cobb-Douglas function of total land and labor:[9]

$$Y_{\mathrm{A}kt} = a_{kt} \bar{X}_k^{\delta} L_{kt}^{1-\delta}. \tag{2}$$

Total land $\bar{X}_k$ is fixed, and there are diminishing returns to labor in the aggregate production function because additional farmers pull less fertile land into agricultural use. The average yield in region k, equal to the average yield of crop k, is consequently declining in the share of land in the region devoted to agriculture, s_{kt}, so that

$$\text{yield}_{kt} = \frac{Y_{\mathrm{A}kt}}{X_{kt}} = a_{kt} s_{kt}^{-\delta}. \tag{3}$$

To highlight the effect of HYVs, assume that the crop-specific (and region-specific) productivity parameter a_{kt} can be decomposed as $a_{kt} = a_{kt}^{\mathrm{HYV}} \bar{a}_k z_t u_{kt}$, where $a_{kt}^{\mathrm{HYV}} = 1$ if no HYVs are available and $a_{kt}^{\mathrm{HYV}} > 1$ if HYVs are grown. The magnitude of a_{kt}^{HYV} depends on how much the available varieties improve yields and on the uptake of HYVs. The parameter $\bar{a}_k$ is a crop-specific time-invariant productivity level, z_t is a country-wide productivity trend, and u_{kt} is an idiosyncratic productivity shock to crop k with mean 1. We can consequently write

[9] This is an application of Houthakker (1955), who demonstrated that Leontief production functions at the establishment level and Pareto-distributed productivity levels imply that the aggregate production function is Cobb-Douglas. See eq. (A2) for the intermediate steps needed to derive this equation.

$$\ln \text{yield}_{kt} = \ln a_{kt}^{\text{HYV}} - \delta \ln s_{kt} + \ln \bar{a}_k + \ln z_t + \ln u_{kt}. \tag{4}$$

In our empirical analysis, we estimate $\ln a_{kt}^{\text{HYV}}$ from a regression analogue to this equation.

We now turn to the aggregate implications of introducing HYVs. To simplify notation, we abstract from factors other than HYVs affecting yields by setting $\ln z_t = \ln u_{kt} = 0$. Let $t = 0$ be the period preceding the Green Revolution. The direct contribution of HYVs to aggregate yields, keeping allocations of land and labor constant at their pre–Green Revolution levels, is

$$\text{GR}_t = \sum_{k=1}^{N} a_{kt}^{\text{HYV}} \frac{Y_{\text{A}k0}}{Y_{\text{A}0}}, \tag{5}$$

where $Y_{\text{A}k0}/Y_{\text{A}0}$ is the share of crop k in aggregate crop production before the Green Revolution. In our empirical analysis, we use equation (5) and our estimates of $\ln a_{kt}^{\text{HYV}}$ to construct an empirical counterpart of GR_t, which allows us to identify country-level effects of the Green Revolution. A GR_t value of 1.5 implies that HYVs, everything else being equal, have increased aggregate yields by 50%. Everything is not equal, however, as the Green Revolution also affected allocations of land and labor. Let period 0 be the last period preceding the Green Revolution. Changes in aggregate yields between period 0 and any period t after the Green Revolution can, under the assumption of free mobility of labor across regions, be written as

$$\Delta \ln \text{yield}_t = \ln \text{GR}_t + \ln \Phi_t - \delta \Delta \ln X_t, \tag{6}$$

where $\Phi_t \equiv (\Sigma_{k=1}^{N}(a_{kt}^{\text{HYV}}\bar{a}_k)^{1/\delta}\bar{X}_k)^{\delta}/(\Sigma_{k=1}^{N} a_{kt}^{\text{HYV}}\bar{a}_k^{1/\delta}\bar{X}_k)$. This equation shows that the effect on aggregate yields from the introduction of HYVs can be decomposed into a direct contribution to the yields of individual crops ($\ln \text{GR}_t$), a productivity gain from reallocation of cropland toward regions growing HYVs ($\ln \Phi_t$), and changes to the extent of cropland ($\delta \Delta \ln X_t$), which affect yields because of decreasing returns to scale. The first two effects are unambiguously positive, whereas the last may contribute positively or negatively, depending on how the Green Revolution affected land use at the extensive margin.

In relatively open economies, higher yields would unambiguously lead to increased specialization in agriculture and consequently increased land use. In relatively closed economies, however, the effect will depend on demand elasticities for food. As a simple model, assume that demand is perfectly inelastic, with individuals consuming a subsistence level of food, denoted $\bar{c}_{\text{A}}$, necessary for survival. Beyond that subsistence level, no more food is demanded. Assume, moreover, that the entire population is in the labor force and that food cannot be stored between periods.

In a closed economy, total demand for food, $\bar{c}_A L_t$, must equal total supply, given by $\text{yield}_t \times X_t$. This equilibrium condition allows us to write yield growth as

$$\Delta \ln \text{yield}_t = \frac{1}{1-\delta} \ln \text{GR}_t + \frac{1}{1-\delta} \ln \Phi_t - \frac{\delta}{1-\delta} \Delta \ln L_t, \qquad (7)$$

which shows why demography is important when evaluating the effect of the Green Revolution. A shock to agricultural productivity will immediately increase yields, but whether the increase can be sustained in the long run depends on the demographic response. If higher yields increase population growth, then $\Delta \ln L_t$ would increase as t increases, putting downward pressure on yields until they are back at their initial level. We do not model fertility and mortality explicitly here, but the literature suggests that a Malthusian effect is far from certain. Higher incomes may, for instance, make parents substitute child quantity for child quality, leading to lower fertility and better education outcomes, as in Becker, Murphy, and Tamura (1990) and Galor and Weil (2000).

Our discussion has so far dealt only with the agricultural sector. To see what the Green Revolution means for the aggregate economy, suppose now that in addition to agriculture, the economy has a manufacturing sector producing output $Y_{Mt} = mL_{Mt}$, where m is a constant productivity term. Labor must be employed in either of the two sectors, meaning that $L_{At} + L_{Mt} = L_t$. GDP per capita can consequently be written as

$$y_t = \frac{Y_t}{L_t} = \frac{Y_A}{L} + \frac{Y_M}{L} = \bar{c}_A + p_t m \frac{L_{Mt}}{L_t}, \qquad (8)$$

where p_t is the relative price of manufactured goods. This equation shows that structural transformation will necessarily accompany any increase in GDP per capita, and our assumptions mean that structural transformation ultimately is driven by yields (because $\Delta \ln (L_{At}/L_t) = -\Delta \ln \text{yield}_t$). For the purpose of our empirical application, it is convenient to focus on growth in GDP per capita (in constant prices), which can be approximated as

$$\Delta \ln y_t \approx \frac{p_0 m}{y_0} \frac{L_{A0}}{L_0} \Delta \ln \text{yield}_t. \qquad (9)$$

A shock to yields is moderated not only by the size of the labor force initially employed in agriculture but also by the term $p_0 m / y_0$, which measures how much more productive workers are in manufacturing than in other sectors. If the manufacturing sector is more productive than agriculture, as is typical in low-income economies (Gollin, Lagakos, and Waugh 2014), then $p_0 m / y_0 > 1$, and reallocation of labor to manufacturing would increase GDP. Equation (9) consequently shows that reallocation of labor to manufacturing might amplify yield growth to such an extent that the effect on GDP per capita is larger than the isolated effect in agriculture.

IV. Research Design

Our empirical analysis follows the same steps as the theoretical framework. We first estimate the effects of the Green Revolution on the relative yields of individual crops and then proceed to study country-level outcomes. Our empirical strategy in both cases relies on the release dates of HYVs, reported in table 1; as noted above, we argue that these release dates are exogenous to individual countries. At the crop level, we use the different release dates and the staggered adoption (or rollout) of treatment to estimate the effect of HYV releases on the yields of affected crops, relative to unaffected crops.[10] Our crop-level estimates of yield gains from HYVs, combined with initial production shares of different crops, allow us to construct an empirical counterpart of GR_{it}, as given in equation (5). The resulting variable is equivalent to a shift-share (Bartik) instrument, which we use for the purpose of identification at the aggregate level (although in reduced form). Many shift-share strategies rely on shift variables that are endogenous at the aggregate level but are assumed to be exogenous to local conditions. The aggregate endogeneity of the shift variable hinders causal inference beyond the local effects. By contrast, we use an exogenous shift variable obtained from our causal crop-level estimates. This allows us to use our estimates to calculate the total contribution of HYVs to economic growth in developing countries.

A. Crop-Level Framework

We estimate the effect of HYVs on crop yields using annual data at the country-crop level. We start by estimating the following event-study version of equation (4):

$$\ln \text{yield}_{kit} = \sum_k \sum_{j \in T_k} \beta_{kj} \cdot \mathbf{1}_{kt}^{t=\tau_k+j} + \delta \ln \text{harea}_{kit} + \mu_{ik} + \mu_{it} + \varepsilon_{kit}, \quad (10)$$

where k indexes crop, i indexes country, and t indexes time. The two terms μ_{ik} and μ_{it} denote, respectively, country-by-crop fixed effects and country-by-year fixed effects, meaning that only within-country time variation in relative yields remains. The country-by-year fixed effects control

[10] Staggered-adoption (rollout) designs have been used to study trade (Autor 2003), health and development (Duflo et al. 2015; Alsan and Goldin 2019), human capital (Acemoglu and Angrist 2000), and natural resource extraction (Bartik et al. 2019). Shift-share instruments have likewise been applied in many settings, including immigration (Card 2001), trade (Autor 2003), health and economic growth (Acemoglu and Johnson 2007), and banking (Greenstone, Mas, and Nguyen 2020).

for all country-specific time variation affecting all crops, including weather shocks and trends toward intensification and mechanization of the agricultural sector.

We expect the introduction of HYVs of crop k to increase their yields relative to other crops in a given country. To capture this effect in the regression, we include as explanatory variable an indicator function $\mathbf{1}_{kt}^{t=\tau_k+j}$ that takes a value of one j years after the global release year of the first HYV of crop k, which we denote τ_k. As mentioned in section II, τ_k varies across crops, providing us with exogenous time variation. The baseline regression includes pure-control crops for which no HYVs were introduced and for which the indicator takes the value zero for the entire period. The error term, ε_{kit}, captures country-specific trends in relative yields, so the coefficient β_{kj} measures by how much the relative yield of crop k in the average country has changed j years after the introduction of HYVs relative to a benchmark year. A natural benchmark is the year before introduction of HYVs, so we define $T_k = \{-10, .., -2, 0, 1, .., 2010 - \tau_k\}$. If HYVs provided the only global shock to relative crop yields at the specified release dates and our identifying assumptions are otherwise correct, then the estimated β_{kj} would be the empirical counterpart of $\ln a_{kt}^{\mathrm{HYV}}$ in the theoretical framework outlined in section III (see also app. C). This provides us with the testable hypothesis that $\beta_{kj} > 0$ after the release of the first HYV of crop k (i.e., for $j \geq 0$) and $\beta_{kj} = 0$ before (i.e., for $j < 0$). We control for harvested area, $\ln \mathrm{harea}_{kit}$, to map the estimating equation into our theoretical framework. We thereby take into account that higher yields of crop k will lead its production to expand into less suitable areas, meaning that the estimated β_{kj} should be interpreted as the effect on relative yields for a fixed allocation of land.[11] Equation (10) also allows for differential effects of HYVs across crops, a reasonable assumption, given botanical differences and differences in research intensity. However, for expositional reasons, we also estimate a version of the event study in which we estimate the average effect of HYVs across all treated crops by imposing $\beta_k = \beta$ and a version in which we estimate separate average effects for different crop types (i.e., cereals, pulses, and roots and tubers).

In addition to our event studies, we estimate the effect of HYVs on crop yields using a difference-in-difference strategy with staggered variation in the timing of treatment coming from the HYV release dates. Our difference-in-difference estimating equation is

$$\ln \mathrm{yield}_{kit} = \sum_k \alpha_k \cdot \mathbf{1}_{kt}^{\mathrm{HYV}} \cdot t + \delta \ln \mathrm{harea}_{kit} + \mu_{ki} + \mu_{it} + \varepsilon_{kit}, \qquad (11)$$

[11] Harvested area is obviously an endogenous control, but excluding it from the regression has quantitatively unimportant effects on our baseline estimates.

where $\mathbf{1}_{kt}^{\mathrm{HYV}}$ is an indicator equal to one in years after the release of the first HYV of a crop k. For pure-control crops, $\mathbf{1}_{kt}^{\mathrm{HYV}}$ is zero throughout the sample period. Because the indicator is interacted with a linear year trend, we assume a trend break rather than a mean shift in yields of a crop after the first HYV release. This assumption is not taken for granted: trend breaks are clearly visible in our event studies below. A priori, however, we would also expect to see such a pattern. Adoption of HYVs happens gradually, and aggregate yields follow a trend closely linked to the adoption rate, even if adoption of HYVs at the farm level causes an immediate jump in yield levels.[12] Moreover, breeding did not stop with the first HYV released for a given crop, and newer vintages of HYVs often perform better in terms of yields, disease resistance, and drought tolerance. How quickly yields increase after the first release of a HYV variety of crop k is captured by the coefficient α_k. As with the event studies, we estimate versions of equation (11) where we impose common trends of treated crops, that is, $\alpha_k = \alpha$, and common trends within crop groups, that is, $\alpha_k = \alpha_\kappa$ where $\kappa =$ (cereals, pulses, roots and tubers).

The main identifying assumption of our difference-in-difference strategy is that if HYVs of crop k had not been released, yields of crop k would have followed the same trend as yields of crops with no HYV releases. While this counterfactual is unobservable, it is supported by our event studies, which show that yields of crop k followed the same trend as yields of other crops before the first HYV of crop k was released.

B. Country-Level Framework

From our crop-level estimates, $\hat{\alpha}_k$, we obtain an empirical counterpart of equation (5),

$$\widehat{\mathrm{GR}}_{it} = \sum_{k=1}^{N} \exp(\hat{\alpha}_k \mathbf{1}_{kt}^{\mathrm{HYV}} t) \frac{Y_{\mathrm{A}ik0}}{Y_{\mathrm{A}i0}}, \tag{12}$$

where the observed pre–Green Revolution production shares, $Y_{\mathrm{A}k0}/Y_{\mathrm{A}0}$, are measured in constant prices and averaged over 1961–64 to reduce noise from, for example, weather shocks. As above, the indicator $\mathbf{1}_{kt}^{\mathrm{HYV}}$ is zero the entire period for crops of which HYVs have not been developed. The log of $\widehat{\mathrm{GR}}_{it}$ is analogous to a shift-share instrument for log yields. Because $\ln \widehat{\mathrm{GR}}_{it}$, by construction, is zero before the onset of the Green Revolution, we can interpret $\ln \widehat{\mathrm{GR}}_{it}$ in any year after that as the (approximate) predicted exogenous growth contribution of HYVs to aggregate yields under the assumption of fixed allocations of land and labor. For shorthand, we therefore refer to $\ln \widehat{\mathrm{GR}}_{it}$ as "predicted GR yields." To

[12] A seminal paper by Griliches (1957) documents this process in the case of hybrid maize in the United States. For evidence of the gradual adoption of HYVs, see fig. A6*F*.

estimate the general equilibrium effect on aggregate yields (no longer keeping allocations of land and labor fixed) and other agricultural outcomes, we run regressions of the form

$$\ln y_{it} = \lambda \ln \widehat{\mathrm{GR}}_{it} + \mathbf{X}_{it}\rho + \mu_i + \mu_t + \varepsilon_{it}, \tag{13}$$

where y_{it} is the outcome of interest in country i at year t, $\mathbf{X}_t$ is a vector of control variables, μ_i and μ_t are country and year fixed effects, respectively, and ε_{it} is an error term. Under the model in section III, we should expect the coefficient λ to be a composite of the direct effect of HYVs for fixed allocations of land and labor, the effect of reallocation of these inputs, and a demographic response to higher yields (see eq. [7]). We bootstrap the standard errors in the regression to take into account that the predicted vector of GR yields is a generated regressor. Specifically, we use a version of the wild-cluster restricted bootstrap procedure proposed by Cameron, Gelbach, and Miller (2008), adapted to our setting.[13]

We also estimate the effect of the Green Revolution on economy-wide outcomes such as GDP per capita and population size. Equation (9) in our theoretical framework shows that we should expect yields to have larger effects on GDP per capita in countries with high initial agricultural employment shares, and it seems reasonable to assume that the same is true for other economy-wide variables. Our baseline estimates of the economy-wide effects of the Green Revolution are consequently obtained from the following regression:

$$\ln y_{it} = \lambda(\ln \widehat{\mathrm{GR}}_{it}) \times \frac{L_{\mathrm{A}i0}}{L_{i0}} + \mathbf{X}_{it}\rho + \mu_i + \mu_t + \varepsilon_{it}, \tag{14}$$

where $L_{\mathrm{A}i0}/L_{i0}$ is the observed initial employment share in agriculture and the remaining variables are as defined above. Because our main explanatory variable now is an interaction, we include in the control set $\mathbf{X}_{it}$ the initial agricultural employment share fully interacted with year fixed effects. Again, we adjust the standard errors to take the generated regressor into account.

Our theoretical framework also predicts that the magnitude of the yield effect on GDP per capita will depend on the productivity level in nonagriculture relative to the aggregate productivity level ($p_0 m/y_0$ in

[13] Our implementation of the bootstrap resembles the procedure for 2SLS (two-stage least squares) regressions with clustered standard errors described by Roodman et al. (2019). What is different in our setting is that our baseline regressions are reduced form, rather than 2SLS, and that our reduced-form generated regressor is aggregated from first-step estimates using eq. (12). We report bootstrapped standard errors based on 1,000 replications. Both Cameron, Gelbach, and Miller (2008) and Roodman et al. (2019) recommend inference based on bootstrapped p-values, which in some cases has slightly better asymptotic properties than inference based on bootstrapped standard errors, but the two methods result in indistinguishable levels of statistical significance in our application.

the model). We do not observe this quantity in our data, but treatment heterogeneity in this dimension will to some extent be reflected in our estimated λ. One reason is that a lower agricultural employment share will give nonagricultural productivity a greater weight in aggregate productivity and thereby push $p_0 m/y_0$ toward unity. Another reason is, as shown by Gollin, Lagakos, and Waugh (2014), among others, that the agricultural productivity gap is declining in the level of development, which is reasonably well approximated by the agricultural employment share.

Equation (12) demonstrates that our identifying variation comes from the mix of crops that countries were producing before the Green Revolution. Another way of specifying this variation is, as Goldsmith-Pinkham, Sorkin, and Swift (2020) argue in their discussion of shift-share instruments, to use the initial production shares $Y_{\mathrm{A}ik0}/Y_{\mathrm{A}i0}$ as separate instruments. The relevance of each instrument is essentially our estimated effect of HYVs on crop yields, that is, $\hat{\alpha}_k$, making our identification strategy equivalent to a difference-in-difference strategy with a continuous treatment intensity measured by $Y_{\mathrm{A}ik0}/Y_{\mathrm{A}i0}$. We can use this equivalence to test whether our outcome variables in countries highly exposed to HYVs followed the same trend as in other countries before the Green Revolution. Absence of such pretrends would lend credibility to our main identifying assumption that countries exposed to HYVs would have followed the same trajectory as countries less exposed to HYVs if the Green Revolution had not happened. To make this idea operational in an event study, we exploit that HYVs of maize, rice, and wheat were released almost simultaneously in the mid-1960s, whereas HYVs of other crops started to emerge only around 1980. We can therefore investigate what happened to countries initially specialized in maize, rice, and wheat around the onset of the Green Revolution. For this analysis, we use the event-study equation

$$\ln y_{it} = \sum_{j=1950}^{2010} \gamma_j \Omega_i \times I_t^j + \mathbf{X}_{it}\rho + \mu_i + \mu_t + \varepsilon_{it}, \tag{15}$$

where Ω_i is the sum of the initial production shares of wheat, rice, and maize. Because Ω_i is observed, we do not need to bootstrap standard errors in the event studies. We interact Ω_i with a full set of year fixed effects (I_t^j), where the omitted year of comparison is 1964. It would support our empirical approach if we find no evidence of trends in the outcomes related to Ω_i before the onset of the Green Revolution in 1965. And if HYVs affected country-wide outcomes, we should see countries growing wheat, rice, and maize diverge from other countries starting from 1965 and at least up to the 1980s, when HYVs of other crops started to diffuse. From that point onward, the relative performance of countries growing the early HYV crops depends on the effect of HYVs on their yields compared to late HYV crops; that is, on the parameters α_k.

V. Data

The key variable in our analysis is crop yield, defined as physical units of crop production per harvested area. Starting from 1961, FAO reports data on these variables on an annual basis for 158 different crops in all UN member countries. The first HYVs of wheat, rice, and maize became available in the mid-1960s, so to test for pre-HYV trends, we collected historical data from other sources to supplement the FAO data. Our main source for pre-1961 data on harvested area and crop production is Mitchell (1982). We obtain additional data from Rose (1985), the *Yearbook of the League of Nations* (various issues), the *Statistical Yearbook of the United Nations* (various issues), and reports from the Economic Research Service of the US Department of Agriculture and the national statistical agencies in China and India. In total, we have collected about 13,500 observations of annual crop production and harvested area in developing countries covering the period 1920–61. Our historical data set spans fewer countries and crops than the FAO data but is quite comprehensive for wheat, rice, and maize—the crops for which we need historical data to test for pre-HYV trends. Other major crops, such as barley and cassava, also have adequate coverage. To aggregate yields across crops, we switch from physical units to value units, using international farm-gate prices in 1966 as our price weights.[14]

In the crop-level analysis, we study 16 mandate crops of the international agricultural research centers (IARCs), as listed in table 1.[15] We compare the treated crops to each other and to a set of botanically similar pure-control crops for which no breeding of HYVs has taken place

[14] Ideally, we would use country-level prices for all crops before the Green Revolution. Comprehensive farm-gate price data in local currency are available from FAO, but the data start in 1966, cover only a subset of countries, and have many missing observations. The data do, however, contain sufficient observations to compute the average relative price of each crop in a common currency, giving us the estimate of international farm-gate prices in 1966 that we use in the aggregation (see app. D.1.3). The year 1966 is sufficiently early in the Green Revolution that HYVs had not increased yields to the extent that relative world prices were affected. In a robustness check, we show that we obtain similar conclusions if we aggregate by nutritional value, an approach to aggregation also used by Galor and Özak (2016). Nutritional values by crop are reported in table A2. See fig. A5 for a comparison with prices.

[15] There are 19 IARC mandate crops, but we exclude plantains, potatoes, and soybeans from the sample for the following reasons. Plantains are botanically sterile, making breeding problematic; only in recent years, using genetic engineering tools, has varietal improvement become a real possibility for research. For potatoes, we cannot establish a firm cutoff date from which the first HYVs became available, because improved varieties from North America and Europe were in wide use in developing countries before the initiation of the IARCs. For soybeans, the role of the IARCs is small compared to that of the private sector, which markets GM varieties in many countries; this could lead to an upward bias in our estimates. We note that the private sector is also active in maize breeding, but in most developing countries, the private sector works with parent material that comes from the IARCs. The private sector is not a particularly major presence in varietal research in developing countries for the other crops that we consider in this paper.

(see sec. IV.A). According to the FAO classification, all the treated crops belong to one of the following three crop types: cereals, pulses, and roots and tubers. For this reason, we choose as pure-control crops all varieties of cereals, pulses, and roots and tubers for which neither the IARCs, nor public sector researchers in developed countries, nor commercial breeders have exerted significant research effort into developing HYVs. Nine crops (predictably minor ones) fulfill this condition: bambara beans, buckwheat, canary seed, fonio, lupins, quinoa, taro, vetches, and yautia.

Because we model the effects of the Green Revolution via its impact on food production, our country-level analysis correspondingly focuses on yields of food crops—or more specifically, cereals, pulses, and roots and tubers.[16] These are the crops for which the FAO data are most comprehensive and reliable. The 35 crops belonging to these three groups account for about 80% of the total harvested area in our sample of countries. We use all the crops in these categories to calculate the predicted GR yields—not just the treated crops and the pure-control crops. Crops such as oats, for which plant breeding took place outside the auspices of the IARCs, are included in calculating the initial production shares; so are potatoes, which benefited from breeding outside the IARCs as well as from the International Potato Center. However, we emphasize that scientific advances in these crops do not contribute to our predicted GR yields.

The IARCs targeted developing countries, so all European countries, all former Soviet republics, Australia, Canada, Israel, Japan, New Zealand, and the United States are excluded from the sample. In our baseline sample, we also exclude countries with fewer than 10,000 ha of arable land devoted to food production. We additionally exclude the 10 largest oil producers measured by barrels per capita in 2017 (Brunei, Gabon, Equatorial Guinea, and seven countries in the Middle East) and Botswana, whose diamond production makes it as dependent on natural resource extraction as the excluded oil countries.[17] We end up with a baseline sample of 90 countries for which we have GDP data and a sample of 86 countries for which we have data for agricultural employment.[18]

[16] We exclude fruits, nuts, and vegetables from our baseline crop sample because of the lack of comparability of yield data and because of relatively many missing observations. For many fruit and nut crops, yield is not a particularly meaningful measure. Fruit trees, especially in developing countries, may be planted in the back yard or as isolated trees in a field. In those settings, output per unit land area is not a useful concept. To some extent, the same is true of continuously harvested vegetables that can be grown productively in small spaces. Yield is a better measure for so-called field crops that are harvested fully at a moment in time.

[17] The reason for excluding these countries is that agriculture plays a fundamentally different (and frequently insignificant) economic role in these economies. Agricultural productivity growth does not seem relevant in these countries.

[18] Table A4 shows the specific countries included in our baseline sample, which we consider most appropriate for evaluating the Green Revolution. We show in app. A12 that our results are robust to changing the sample restrictions.

Our crop-level sample period is 1945–2010. The sample is unbalanced, as the pre-1961 data cover fewer crops and countries. The unbalanced nature of the sample does not affect our crop-level analysis, but it makes aggregate variables based on these data fluctuate for purely statistical reasons. Our baseline country-level sample period for agricultural variables is consequently restricted to 1961–2010. When we use GDP per capita or population size as outcomes, we are able to begin the analysis in 1950.

Further details about our crop-level data, as well as the sources for data on GDP per capita, population size, and the other outcome variables in our analysis, can be found in appendix D.

VI. Results

A. Crop-Level Effects of the Green Revolution

The impact of HYVs on crop yields is clearly visible in the event-study graphs in figure 2. Based on equation (10), the graphs show estimated nonparametric trends in relative yields of a crop before and after the first HYV of the crop was released. Figure 2*A* displays estimates under the assumption that HYVs have identical effects on yields of all 16 crops for which HYVs were released. For event years before the first HYV release (the period up until the vertical line), the estimated coefficients are close to zero and statistically insignificant, implying that yields of treated crops and untreated crops followed the same trend before the first HYV release. The absence of differential pretrends supports the identifying assumption in our difference-in-difference estimates below. After the first HYV was released, the relative yield of the treated crop significantly increases.

Converting our estimates from logarithms to levels, we find that relative yields are on average 9% higher 10 years after a HYV release ($\beta_{10} = 0.09$) and 75% higher after 40 years ($\beta_{40} = 0.56$). The gradual increase in yields happens both because adoption is gradual, along an extensive margin, and because successive vintages of HYVs of a crop increase yields beyond what the first HYV could achieve. Our estimated magnitudes are consistent with the micro-level literature, surveyed in Evenson and Gollin (2003b), which shows that HYVs typically have at least 50% higher yields than traditional varieties for a given set of inputs. Inputs are not fixed, however. Many HYVs respond better to fertilizer and other inputs than traditional varieties, raising yields still further; gains of the magnitude observed in figure 2 are not unexpected, in cases when HYV adoption is widespread.

Figures 2*B*–2*D* report separate event studies for cereals, roots and tubers, and pulses, respectively, thereby allowing for yield gains to differ for botanically different crops. For all three crop types, we find no evidence of pretrends and highly significant increases in relative yields following the

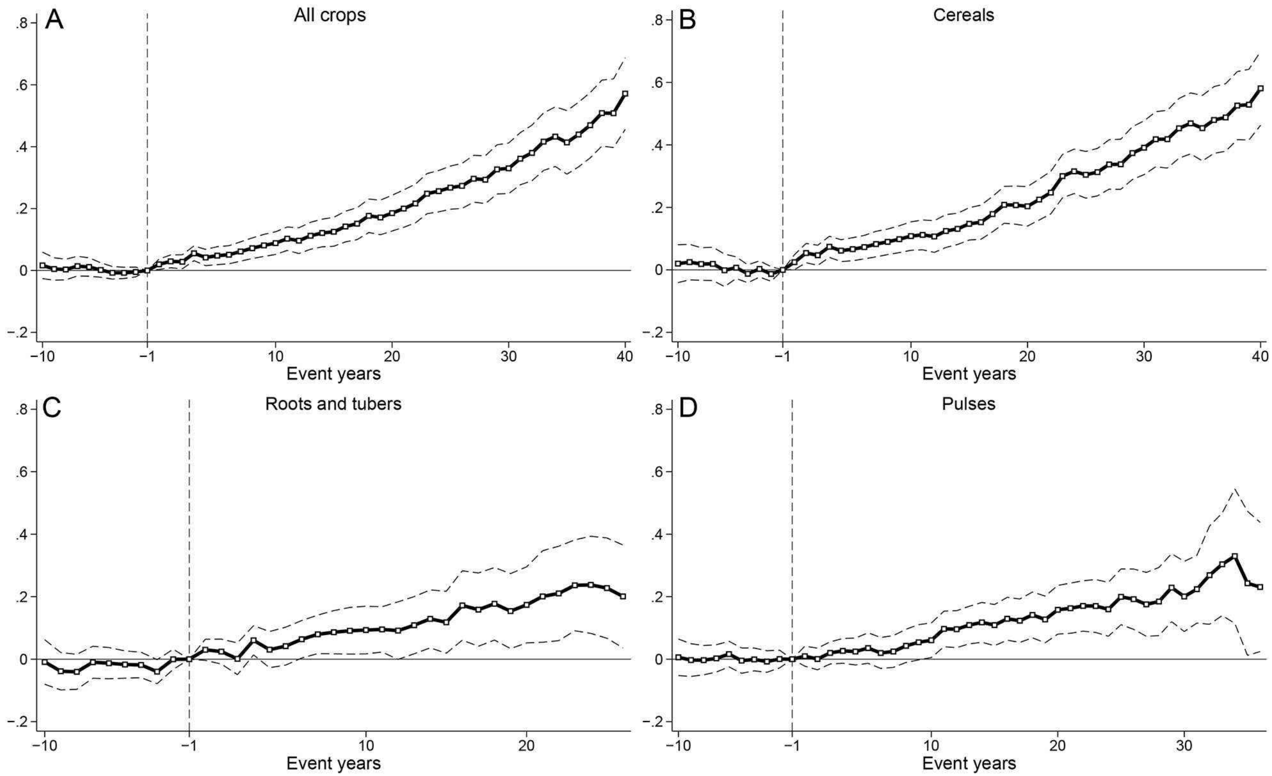

Fig. 2.—Baseline crop-level event-study estimates based on equation (10). The dependent variable is ln yields. The sample period is 1945–2010. In *A*, we show estimates under the assumption that the treatment effect is of the same magnitude across crops. In *B*–*D*, we report effects by crop type (cereals, roots and tubers, pulses), all from the same regression. Both regressions (i.e., *A* and *B*–*D*) include controls for ln harvested area, country-by-year fixed effects, and crop-by-country fixed effects. The omitted comparison event-year is the year before an HYV release (vertical line). The dashed curves indicate the 95% confidence bands. Standard errors are clustered at the crop-by-country level.

first HYV release. The magnitudes differ across the three crop groups, with HYVs of cereals having larger impacts.[19]

Table 2 reports difference-in-difference estimates based on equation (11). The results in columns 1 and 2 correspond to the event-study graphs in figure 2, except that we now replace the nonparametric trends with linear, post-HYV trends. We find positive and significant effects from HYVs of all three crop types, but a comparison of the estimates in Column 2 shows that the yield gain for cereals is about 44% larger than the gain for roots and tubers and 85% larger than the gain for pulses. At the level of individual crops (see table A5), we find positive effects of HYVs on almost all crops, although the coefficients are imprecisely estimated for minor crops with relatively few observations (e.g., faba bean). The largest impacts of HYVs are on the yields of wheat, maize, rice, and barley, but not all cereals have seen similarly large gains after the Green Revolution. Yield gains for sorghum and millet have, for instance, been modest.

Columns 3–7 show the robustness of the results in column 2. In column 3, we exclude yield observations with low data quality.[20] The estimates are slightly higher than our baseline, suggesting that attenuation bias is present but small. In column 4, we disregard the historical data we have collected and use only the FAO data starting in 1961. The result is almost identical to our baseline, so our estimates are not driven by a change in data source. In column 5, we exclude countries hosting an IARC from the sample to show that host countries are not driving our results. In column 6, we exclude the pure-control crops to show that our results are robust to the selection of these crops. In column 7, we control for crop-type linear trends and thereby remove linear pretrends from the estimates. Unsurprisingly, given the picture in figure 2, we find estimates highly similar to those in column 2.

In table A6, we demonstrate that the harvested area of treated crops increases relative to that of untreated crops, meaning that farmers substitute toward crops with HYVs, which is consistent with both our theoretical framework and simple economic logic. We also provide suggestive evidence from a limited sample that the relative price of a crop falls after the first HYV of the crop is released. Our theoretical framework assumes

[19] Figures A1 and A2 report separate event studies for each of the treated crops in the categories cereals, and roots/tubers. Crops in the pulses category are minor crops, with too few observations to reliably estimate the rather flexible event-study specification for them separately.

[20] Our high-quality data set excludes FAO estimates for countries with no official statistics or with official data that are (1) unvarying from year to year, (2) have less than three significant digits (we remove all zeroes trailing the last nonzero digit and count the number of digits left), and (3) are based on crops for which harvested area is less than 1,000 ha. The two first restrictions eliminate data points that are crude estimates, and the last restriction removes outliers from small sample sizes in the underlying surveys.

TABLE 2

Estimated Effect on Crop Yields

	Baseline for All Crops (1)	Baseline by Crop Type (2)	High-Quality Yield Data (3)	FAOSTAT Yield Data (4)	Excluding IARC Countries (5)	Excluding Pure-Control Crops (6)	Crop-Type Linear Trends (7)
HYV × t	.013***						
	(.001)						
HYV × t, cereals		.013***	.015***	.012***	.012***	.014***	.017***
		(.001)	(.002)	(.001)	(.001)	(.002)	(.002)
HYV × t, roots/tubers		.009***	.010***	.009***	.009***	.012***	.010**
		(.003)	(.003)	(.003)	(.003)	(.003)	(.004)
HYV × t, pulses		.007***	.009***	.007***	.006***	.009***	.007***
		(.002)	(.002)	(.002)	(.002)	(.002)	(.002)
ln harvested area	−.023**	−.022**	−.009	−.028**	−.025**	−.021*	−.023**
	(.010)	(.010)	(.016)	(.011)	(.011)	(.011)	(.010)
Country × year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country × crop FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Crop-type linear trends	No	No	No	No	No	No	Yes
Observations	45,184	45,184	27,313	43,256	40,342	41,524	45,184
Countries	90	90	90	90	83	90	90

Note.—Estimates based on eq. (11). The dependent variable is the log value of production per hectare. The explanatory variable is an indicator for the release year of an HYV in a given crop, interacted with a linear time trend (HYV × t) in col. 1. This variable is further interacted with a crop-type indicator (cereals, roots and tubers, pulses) in cols. 2–7. Columns 1 and 2 constitute the baseline specification. In col. 3, we use only high-quality data. In col. 4, only FAOSTAT (FAO Statistical Database) data are used, so the sample period is 1961–2010. In the remaining columns, the sample period is 1945–2010. In col. 5, we exclude from the sample countries hosting an international agricultural research center. In col. 6, pure-control crops are excluded, such that we include only crops for which HYVs became available within the sample period. Standard errors (in parentheses) are clustered at the crop-by-country level. FE = fixed effects.

* $p < .1$.
** $p < .05$.
*** $p < .01$.

that different crops are perfect substitutes and consequently have the same price; but in general, we should expect HYVs to increase the supply of a crop and thereby reduce its relative price, as indeed appears to be the case in the data.

B. The Spatial Distribution of the Green Revolution

Armed with our crop-specific difference-in-difference estimates, we use equation (12) to construct $\widehat{\mathrm{GR}}_{it}$, that is, the estimated exogenous yield shock coming from the Green Revolution calculated under the assumption that the allocation of land and labor is fixed. According to this measure, yields had by 2010, ceteris paribus, increased by 44% in the average country, compared to a counterfactual with no Green Revolution. This average masks substantial spatial variation originating in the mix of crops that countries were growing before the Green Revolution. Countries that devoted much of their crop area to rice and wheat, for which HYVs were released early and had a large impact on yields, benefit relatively more from HYVs in the aggregate. The map in figure 3*A* shows the most important food crop in each country in our sample, measured as the value of production in 1961 (in 1966 prices). Unsurprisingly, Southeast Asia shows up in this figure as rice territory, whereas there is more within-region variation in Africa and Latin America. However, most countries grow a wide range of the 35 crops on which our measure of aggregate food crop yields is based, and the most important crop often accounts for a small fraction of total production. A better way to illustrate the spatial heterogeneity we use as a source of variation in our country-level regressions is to plot $\ln \widehat{\mathrm{GR}}_{it}$ in a map, as we do in figure 3*B*. The map is for 2010, but the predicted GR yields are time varying because of the staggered releases of HYVs of different crop, so plotting it for other years would have resulted in a different picture (see fig. A3 for an illustration of the time variation). Darker shading in the map indicates a higher value of the instrument and consequently a larger growth contribution to yields from HYVs. The map shows that there is substantial variation across Southeast Asia, despite rice being the most important crop in all but a few countries in the region. Still, there is less variation within Southeast Asia than within Africa, where there is more heterogeneity in the crop mix.

C. Aggregate Effects on Agriculture

We now evaluate how the Green Revolution affected agriculture in developing countries by estimating equation (13). Table 3 reports estimates for five different agricultural outcomes. For each outcome, we report our baseline estimate in column 1 and robustness checks in columns 2–5. The

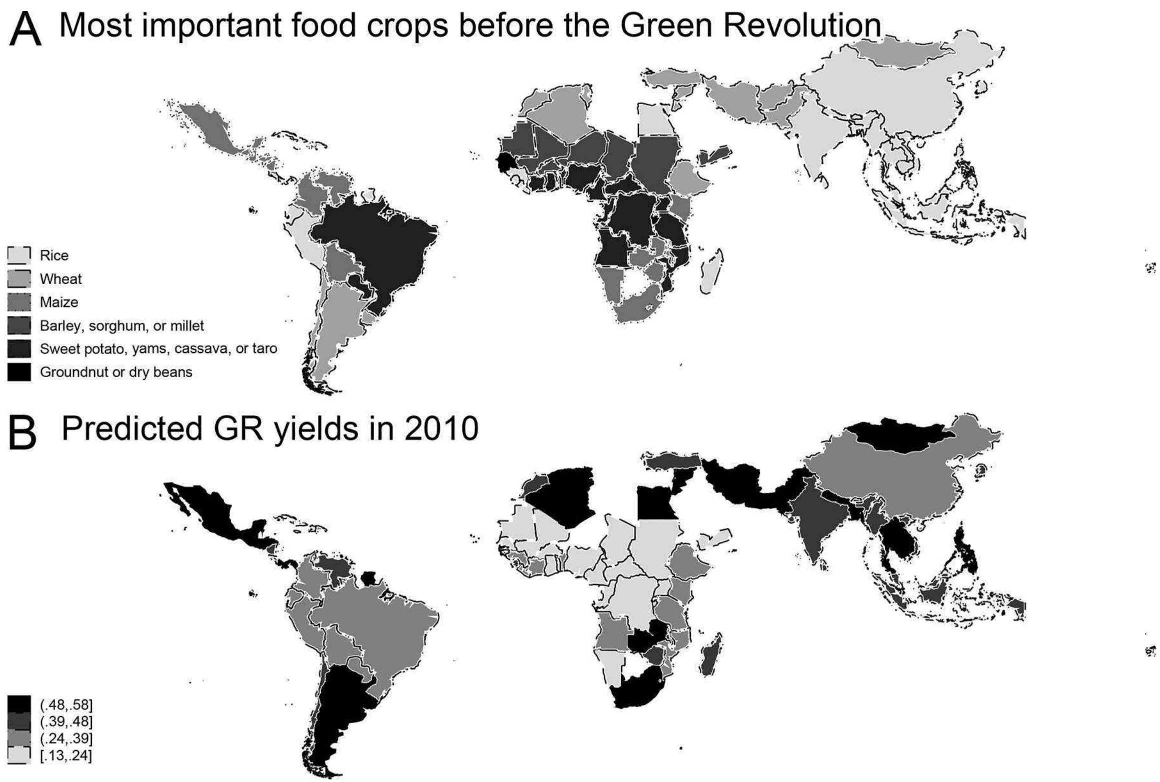

FIG. 3.—The most important food crop for each country, measured at production values in 1966 (*A*), and predicted GR yields in 2010 (*B*), which measures the predicted growth contribution to yields from 1961 to 2010 due to the Green Revolution under the assumption of fixed allocations of land and labor.

TABLE 3

Effects on Country-Level Agricultural Outcomes

	Baseline (1)	Caloric Aggregation (2)	Controlling for Climate (3)	Controlling for Pre-GR Trends (4)	Excluding IARC Host Countries (5)
	A. ln Aggregate Crop Yields				
$\ln \widehat{GR}$	1.73***	1.54***	1.70***	1.72***	1.78***
	(0.41)	(0.41)	(0.33)	(0.42)	(0.43)
Observations	4,500	4,500	4,500	4,500	4,150
Countries	90	90	90	90	83
	B. ln Harvested Area (Food Crops)				
$\ln \widehat{GR}$	−1.90***	−1.78***	−1.91***	−1.88***	−1.95***
	(0.51)	(0.51)	(0.44)	(0.51)	(0.54)
Observations	4,500	4,500	4,500	4,500	4,150
Countries	90	90	90	90	83
	C. ln Harvested Area (Total)				
$\ln \widehat{GR}$	−1.54***	−1.47***	−1.55***	−1.51***	−1.58***
	(0.43)	(0.43)	(0.40)	(0.43)	(0.46)
Observations	4,500	4,500	4,500	4,500	4,150
Countries	90	90	90	90	83
	D. HYV Area Share				
$\ln \widehat{GR}$	1.00***	0.92***	0.99***	1.07***	0.99***
	(0.25)	(0.24)	(0.22)	(0.24)	(0.26)
Observations	3,240	3,240	3,240	3,240	2,960
Countries	81	81	81	81	74
	E. ln Agricultural Employment Share				
$\ln \widehat{GR}$	−1.10***	−1.312***	−1.08***	−1.16***	−1.06***
	(0.37)	(0.35)	(0.33)	(0.37)	(0.38)
Observations	4,300	4,300	4,300	4,300	3,950
Countries	86	86	86	86	79
	F. ln Agricultural Labor Productivity				
$\ln \widehat{GR}$	1.58***	1.89***	1.53***	1.66***	1.60***
	(0.52)	(0.51)	(0.48)	(0.52)	(0.55)
Observations	4,150	4,150	4,150	4,150	3,800
Countries	83	83	83	83	76

Note.—Country-level estimates based on eq. (13). Food crops (in panel B) are cereals, pulses, and roots and tubers. "Total" (in panel C) denotes all crops. HYV area share (in panel D) is the share of agricultural land use for HYVs for 11 major food crops. The value ln Agricultural Labor Productivity (in panel F) is given by logged agricultural value added per agricultural worker. The explanatory variable is $\ln \widehat{GR}$, which measures by how much HYVs have increased yields (under the assumption of fixed allocations of inputs). The sample period is 1961–2010, except for panel D, where the sample ends in 2000. All regressions include country and year fixed effects. Column 1 is the baseline specification. In col. 2, $\ln \widehat{GR}$ is aggregated using caloric content instead of 1966 international prices (as in the baseline). Column 3 controls for climate by including nonparametric controls for temperature and precipitation. Column 4 controls for convergence in the agricultural sector by including ln yields in 1961 (i.e., before the Green Revolution) interacted with year fixed effects. Column 5 excludes countries hosting an International Agricultural Research Center (IARC). Standard errors (in parentheses) are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that $\ln \widehat{GR}$ is generated.

*** $p < .01$.

robustness checks show that for all five outcomes we obtain results similar to the baseline if we aggregate yields using caloric content rather than prices (col. 2), control for weather shocks and climate change (col. 3), control for initial yield levels interacted with year fixed effects (col. 4), or exclude IARC host countries (col. 5).

In the baseline regression in panel A of table 3, we find that the elasticity of aggregate yields with respect to $\widehat{GR}_{it}$ is 1.73, implying that reallocation of land and labor, and possibly additional factor adjustment, amplifies the ceteris paribus effect of HYVs by 73%. The estimated elasticity is significantly larger than unity only at the 10% level, however, so we cannot reject that the magnitude of such amplification is limited. In panels B and C, we find evidence for the Borlaug hypothesis. Panel B shows that total land devoted to food crops fell as a consequence of higher yields (or, perhaps more accurately, rose by less than it otherwise would have done). The effect on total cropland, reported in panel C, is smaller in magnitude, as land devoted to nonfood crops (e.g., cotton, tobacco) was not directly affected by the introduction of HYVs of food crops.

In panel D, we look at the adoption of HYVs, measured by the share of cropland devoted to them. The data, taken from Evenson and Gollin (2003a), end in 2000 and cover only 11 of the HYV crops in our sample (all the major crops are included). Nevertheless, we see a highly significant relationship with our predicted GR yields. The adoption of HYVs is obviously what drives our estimated crop-level yield gain, so this result serves as a consistency check of our identification strategy. Panel E provides evidence that the Green Revolution led to a process of structural transformation. Higher yields freed up labor from agriculture, resulting in a significantly lower agricultural employment share. To the extent that labor productivity is lower in agriculture than in other sectors, this finding amplifies the direct effects of yield gains on total income in the economy. Finally, in panel F, the outcome is a proxy for labor productivity in agriculture, which we calculate by dividing the FAO estimate of the net value of agricultural production (including animal husbandry) by the number of agricultural workers, calculated under the assumption that everyone between 15 and 64 is in employment.[21] We find positive and significant estimates that, in line with the results in panel A, confirm that the Green Revolution has been instrumental for agricultural productivity growth in the developing world.[22]

[21] The FAO data on net production are not compatible with the national accounts data, so we cannot use this measure to construct a meaningful measure of the share of agriculture in GDP.

[22] Event-study graphs for the outcomes in table 3 are reported in fig. A6.

D. *Effects on Economic Development*

We have so far established that HYVs have increased crop yields and fundamentally transformed the agricultural sector since the onset of the Green Revolution. We now turn to the wider implications for economic growth, demography, and development more broadly. We start the analysis by using equation (15) to estimate event studies around the first phase of the Green Revolution, when the first HYVs of wheat, rice, and maize were released almost simultaneously. HYVs of other crops only began to emerge more than a decade later, and with a few exceptions, the yield gains for these latecomers were generally smaller than those for wheat, rice, and maize. Treatment intensity in the event studies is consequently defined as the initial share of wheat, rice, and maize in total food crop production. By estimating such event studies, we are able to detect, and subsequently correct for, possible pretrends that might invalidate our research design.

The event study for GDP per capita in figure 4*A* shows that 10 years after the onset of the Green Revolution in 1965, countries specialized in wheat, rice, and maize begin to have faster income growth than other countries. Before 1965, the growth paths of treated and untreated countries are statistically indistinguishable. One might still worry about the slight positive pretrend in the point estimates, but as shown in figure 4*B*, controlling for pre–Green Revolution GDP growth interacted with year fixed effects eliminates the pretrend while leaving the post–Green Revolution estimates intact. The event-study graph for population growth, in figure 4*C*, shows that countries with higher treatment intensities had faster population growth in the beginning of the sample, but the pattern reverses during the 1980s. While this result may reflect an income effect on fertility choices, it also violates the assumption of parallel trends underlying our difference-in-difference estimates below. Therefore, as with GDP per capita, figure 4*D* reports a version of the event study in which we control for pre–Green Revolution population growth interacted with time fixed effects. This specification is more flexible than adding linear controls for pretrends, as we also allow for mean reversion in the population growth rates, and we are thereby able to control for the possibility that treated and untreated countries were at different stages of the demographic transition before the Green Revolution. As a result, the pretrend is eliminated, and the effect of HYVs on population size becomes slightly stronger.[23]

[23] Figure A7 displays event-study graphs in which we control for pre–Green Revolution linear trends, using the approach in Bhuller et al. (2013) recommended by Goodman-Bacon (2018). We obtain the same conclusion as in fig. 4. Figures A8 and A9 report event-study graphs for wheat, rice, and maize separately. The positive effect on income is visible for all three crops, but the negative effect on population is most pronounced for rice-growing countries. There are no visible pretrends for any of the three crops.

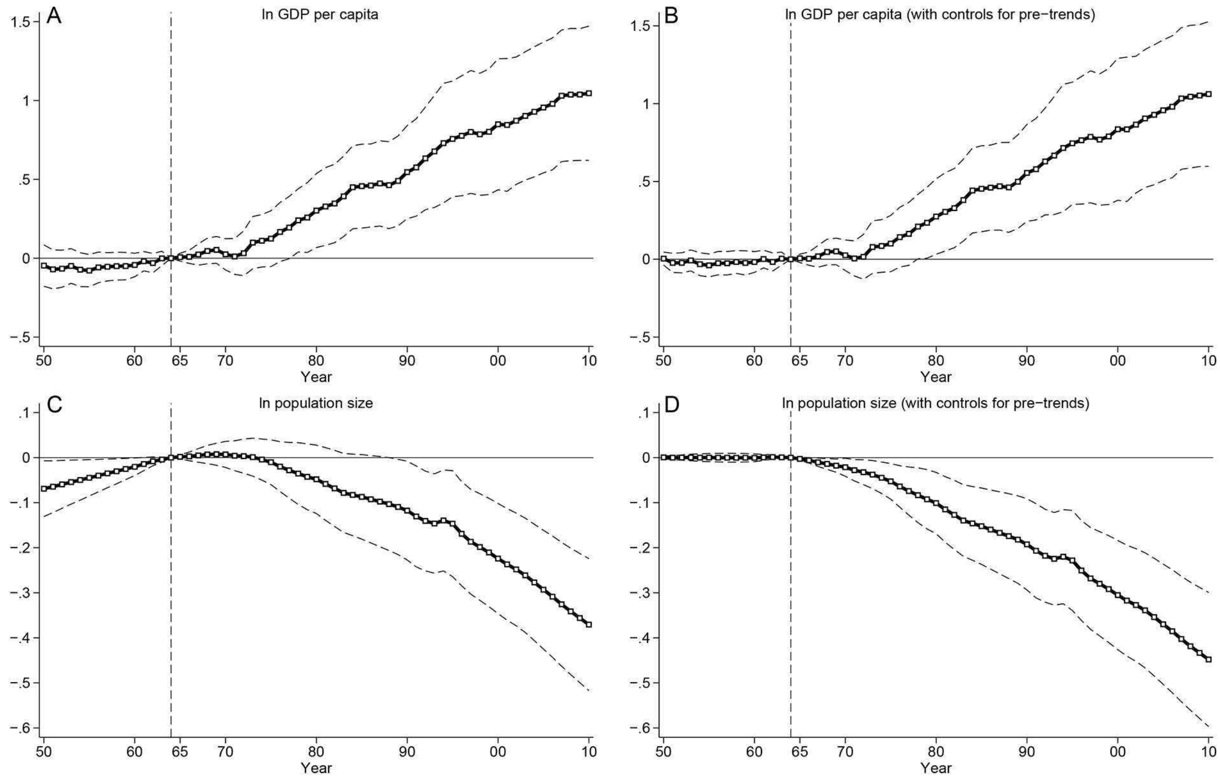

FIG. 4.—Baseline country-level event-study estimates based on equation (15). The explanatory variable (treatment) is the sum of the initial production shares in wheat, rice, and maize interacted with year fixed effects. All regressions control for country and year fixed effects. *B* and *D* additionally control for pre–Green Revolution income and population growth (1950–63) interacted with year fixed effects. The sample period is 1950–2010, and the samples are balanced, with 85 countries. The vertical line indicates 1964, the last pre–Green Revolution period, which is also the omitted comparison year. The dashed curves indicate the 95% confidence bands. Standard errors are clustered at the country level.

Table 4 presents our difference-in-difference estimates of the effect of the Green Revolution based on equation (13). Column 1 corresponds to the event studies in the sense that our predicted GR yield is the main independent variable. The effect of yield growth on the aggregate economy obviously depends on the size of the agricultural sector, which is why, in our baseline regressions, reported in column 2, we interact $\ln\widehat{\text{GR}}_{it}$ with the initial agricultural employment share (L_{Ai0}/L_{i0}). Additionally, we include the initial agricultural employment share interacted with year fixed effects as controls, such that the effect of having a high initial agricultural employment share does not affect our Green Revolution estimates. Implicitly, we also control for the initial stage of development, as GDP per capita and the agricultural employment share are highly correlated.[24] Our baseline estimate for GDP per capita is 2.75 (see table 4, panel A, col. 2). If there were no general equilibrium effects outside agriculture, and if there were no productivity gap between agriculture and nonagriculture, we should expect this estimate to be identical to that for yields, which we found, in table 3, to be 1.73. The larger point estimate is consistent with our results in table 3, showing that higher yields lead to migration of labor out of agriculture toward the more productive nonagricultural sector. A demographic dividend is also part of the story. In panel B of table 4, as well as in the event studies in figure 4, we find a significant negative effect of the Green Revolution on population size. Slower population growth changed the age structure of the population and reduced the dependency ratio, which is why we find, in panel C, a smaller effect of the Green Revolution on GDP per working-age person (defined as people aged 15–64 years) than on GDP per capita. A comparison of the point estimates in panels A and C shows that this demographic divided from the Green Revolution may have accounted for about one-fifth of the total effect on GDP per capita.

Columns 3–6 report selected robustness checks to our baseline results in column 2. In column 3, we follow our approach from the event studies above and control for pre–Green Revolution growth in GDP per capita and population, both interacted with time fixed effects, to eliminate the possible heterogeneity coming from differential pretrends. The other robustness checks mimic those in table 3. The results show that our baseline results are robust to these alternative specifications, including controlling for climate change, which has been found to have a negative impact on economies of the developing world (e.g., Dell, Jones, and Olken 2012 and Burke, Hsiang, and Miguel 2015) and obviously also affects agriculture.

[24] Event-study graphs corresponding to this specification are reported in fig. A10. In addition, marginal plots show that the effects of the Green Revolution on income and population are most pronounced in countries more dependent on agriculture, whereas no such cross-country heterogeneity is present for agricultural outcomes (fig. A11). This finding, which is consistent with the model in sec. III, shows that it is unnecessary to interact with the size of the agricultural sector when studying agricultural outcomes.

TABLE 4
EFFECTS ON INCOME AND POPULATION

	Simple Model (1)	Interaction Model (Baseline) (2)	Controlling for Pre-GR Income and Population Growth (3)	Caloric Aggregation (4)	Controlling for Climate (5)	Excluding IARC Host Countries (6)
	A. Income					
$\ln \widehat{GR}$	2.10***					
	(.56)					
$\ln \widehat{GR} \times$ initial AES		2.75***	2.69***	2.60***	2.71***	2.78***
		(.87)	(.88)	(.85)	(.83)	(.90)
	B. ln Population Size					
$\ln \widehat{GR}$	−.55***					
	(.210)					
$\ln \widehat{GR} \times$ initial AES		−.600**	−.84***	−.53**	−.60**	−.63**
		(.26)	(.23)	(.25)	(.23)	(.27)
	C. ln GDP per Working-Age Population					
$\ln \widehat{GR}$	1.61***					
	(.50)					
$\ln \widehat{GR} \times$ initial AES		2.25***	2.27***	2.09***	2.21***	2.29***
		(.79)	(.81)	(.78)	(.76)	(.83)
	Specifications					
Controls (× year FE):						
Initial AES	No	Yes	Yes	Yes	Yes	Yes
Pre-GR income growth	No	No	Yes	No	No	No
Pre-GR population growth	No	No	Yes	No	No	No
Observations	4,473	4,273	4,050	4,273	4,273	3,923
Countries	90	86	81	86	86	79

NOTE.—Country-level estimates based on eqq. (13) and (14). Working-age population (in panel C) is defined by all people in the age group 16–64. The sample period is 1961–2010. All regressions include country and year fixed effects (FE). Column 1 is the simple model. Column 2 is the baseline model, which interacts $\ln \widehat{GR}$ with the agricultural employment share in 1961 (initial AES), while controlling for initial AES interacted with year fixed effects. The remaining columns are robustness checks: col. 3 controls for pre–Green Revolution (pre-GR) growth in GDP per capita and population growth (1950–63), both interacted with year fixed effects. In col. 4, $\ln \widehat{GR}$ is aggregated using caloric content instead of 1966 international prices (as in the baseline). Column 5 controls for climate by including nonparametric controls for temperature and precipitation. Column 6 excludes countries hosting an International Agricultural Research Center (IARC). Standard errors (in parentheses) are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that $\ln \widehat{GR}$ is generated.

** $p < .05$.

*** $p < .01$.

TABLE 5
EFFECTS ON DEMOGRAPHY AND SCHOOLING

	DEPENDENT VARIABLE					
	ln Adult Mortality Rate (1)	ln Infant Mortality Rate (2)	ln Total Fertility Rate (3)	Rate of Natural Increase (4)	Net Migration Rate (5)	Years of Schooling, Age 15–20 (6)
$\ln \widehat{GR} \times$ initial AES	−1.03*	−2.12***	−2.21***	−.036***	−.003	4.07*
	(.56)	(.59)	(.46)	(.011)	(.008)	(2.32)
Controls (× year FE):						
Initial AES	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,300	4,072	4,300	4,300	4,270	3,750
Countries	86	86	86	86	86	75

NOTE.—Country-level estimates based on eq. (14). The rate of natural increase (col. 4) is calculated as the crude birth rate minus the crude death rate. The net migration rate (col. 5) is calculated as the population growth rate minus the rate of natural increase. All regressions include country and year fixed effects (FE). The model is our baseline model, which interacts $\ln \widehat{GR}$ with the agricultural employment share in 1961 (initial AES) and control for initial AES interacted with year. Standard errors (in parentheses) are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that $\ln \widehat{GR}$ is generated.
* $p < .1$.
*** $p < .01$.

We do, however, find slightly larger effects on population size when pretrends are controlled for. Further robustness checks are reported in appendix H.

The demographic response to the Green Revolution is clearly of first-order importance for explaining the income effects, so, in table 5, we investigate the effects on the demographic variables underlying the overall population response.[25] Column 1 reveals a negative and statistically significant effect on adult mortality, with a point estimate of −1.03. The estimate implies that a 1% increase in yields in a country with half the population employed in agriculture would cause adult mortality to decline by half a percentage point (recall that treatment is $\ln \widehat{GR}_{it} \times L_{Ai0}/L_{i0}$). The effect on infant mortality is larger, with a point estimate of −2.12 (see col. 2). Lower mortality would, by itself, increase the population size, so the impact on fertility (or migration) must be negative and large to rationalize our negative population effect. Column 3 shows a negative impact on the total fertility rate. That the magnitude of this effect is larger than the effect on mortality is shown in column 4, where we find a statistically significant negative effect on the rate of natural increase (i.e., natural population growth). The coefficient is −0.036. Column 5 shows that net migration rates were unaffected by the Green Revolution, so the effect

[25] Figures A13 and A14 display the event-study graphs for all the outcomes in table 5.

on population is entirely driven by changes in mortality and fertility. The last column shows positive effects on human capital, consistent with a quantity-quality trade-off. In table A13, we extend the analysis to a range of additional indicators of development.

VII. Lessons and Perspectives

Our analysis shows that HYVs, originating in international research centers, increased the yields of food crops and per capita income in developing countries. By combining these results, we can estimate the total economic impact of the Green Revolution—an economic return to the crop-breeding efforts that took place in the international agricultural research centers. We provide three such estimates, based on different counterfactual scenarios. The first is based on the impact of the Green Revolution in 2010, compared to a counterfactual in which it never happened. The counterfactual scenario is an out-of-sample prediction of our empirical model, as all countries in our sample by 2010 were affected by the Green Revolution. Additionally, we implicitly assume that no alternative sources of growth would have emerged if the Green Revolution had not happened. While useful as a benchmark, this scenario almost certainly overestimates the return to agricultural research at the IARCs. The research breakthroughs of the IARCs might eventually have been achieved by commercial breeders or national research institutions, with diffusion to the developing world still taking place—but later and more slowly. We cannot know how much longer it would have taken, but it is probably not unreasonable to think that the Green Revolution would have been delayed by at least a decade and possibly substantially longer. We therefore construct two alternative scenarios based on the assumption that the Green Revolution was delayed by 10 and 25 years. In all three scenarios, we transform our estimates from logs to levels and aggregate such that the effect sizes we report apply to the developing world as a whole. The results are summarized in table 6.[26]

Our baseline estimates imply that aggregate food crop yields for our sample of countries would have been 49% lower in 2010 had the Green Revolution never happened. They would still have been higher than those in 1964, but the Green Revolution has accounted for as much as three-quarters of yield growth since then. We find similarly large effects for GDP per capita, which would have been 51% lower in the counterfactual scenario. Taken at face value, this estimate means that the Green Revolution has been responsible for about half of total growth in GDP per capita in our sample period. The average annual population growth rate would

[26] A detailed description of these calculations can be found in app. I.

TABLE 6
COUNTERFACTUAL SCENARIOS

	Levels				Loss/Gain Compared to Actual (%)		
	Actual (1)	Delayed 10 Years (2)	Delayed 25 Years (3)	No GR (4)	Delayed 10 Years (5)	Delayed 25 Years (6)	No GR (7)
Yield in 2010 (2010 USD/ha)	812	687	541	416	−15	−33	−49
Yield growth/year 1965–2010 (%)	2.0	1.6	1.1	.5	−19	−46	−75
GDP/capita in 2010 (2011 USD)	7,580	6,306	4,885	3,693	−17	−36	−51
In LDCs as of today	2,051	1,659	1,223	922	−19	−40	−55
In low–GR effect tercile	6,618	5,926	5,075	4,539	−10	−23	−31
In high–GR effect tercile	5,295	4,293	3,162	2,221	−19	−40	−58
In low–initial income tercile	6,076	4,992	3,826	2,841	−18	−37	−53
In high–initial income tercile	13,881	11,873	9,475	7,364	−14	−32	−47
GDP/capita growth/year 1965–2010 (%)	3.0	2.6	2.0	1.4	−14	−33	−54
Population in 2010 (millions)	5,296	5,519	5,850	6,231	4	10	18
Population growth/year 1965–2010 (%)	1.9	2.0	2.2	2.3	5	11	19
Population growth 1965–2010 (millions)	3,105	3,329	3,659	4,040	7	18	30
GDP in 2010 (trillion 2011 USD)	40.1	34.8	28.6	23.0	−13	−29	−43
Cumulative GDP 1965–2010 (trillion 2011 USD)	724	641	557	515	−12	−23	−29

NOTE.—Effect of the Green Revolution (GR) for the developing world as a whole implied by our estimates. We define the developing world as the 83 countries in our sample for which we have data on all required variables. We compare actual values in 2010 to three counterfactuals: (1) a 10-year delay in arrival of the Green Revolution, (2) a 25-year delay in the arrival of the Green Revolution, and (3) a scenario with no Green Revolution. LDCs are the least developed countries, as defined by the United Nations. The low (high)–GR effect tercile subsample is the third of the countries in the full sample with the lowest (highest) impact of the Green Revolution measured by predicted GR yields multiplied by the agricultural employment share. The low (high)–initial income tercile is the tercile with the lowest (highest) GDP per capita in 1964. Dollar values are PPP adjusted, except for yields, which are aggregated using 1966 world prices and converted into 2010 dollars.

have been 2.3% in the period 1964–2010 without the Green Revolution, 0.4 percentage points higher than it actually was.

Turning to the more plausible second counterfactual scenario, we find that a 10-year delay in the onset of the Green Revolution would have cost the entire developing world, as one would have defined it in 1960, a per capita loss of US$1,273 (PPP adjusted) in 2010, corresponding to 17% of GDP per capita. The dollar amount is large, in part because our sample includes countries such as Chile and South Korea, which grew relatively rich during the period we study. The comparable amount for today's least developed countries is US$392. The population of the developing world would have been about 4% higher, corresponding to 223 million people, had the Green Revolution been delayed by 10 years. By combining this estimate with our estimate for GDP per capita, we find that HYVs developed at the IARCs added roughly US$5 trillion to total GDP in the developing world in 2010 alone, and a cumulative US$83 trillion to GDP since the IARCs were founded. To benchmark this amount, it represents approximately one year of global GDP in 2021. The cumulative GDP gain is, unsurprisingly, substantially larger if we assume that the Green Revolution would have been delayed 25 years rather than just 10 years in the absence of the IARCs. We do not have an exact estimate of how much money the IARCs have spent on developing new HYVs, but by any plausible estimate, the return on investments in the IARCs has been remarkable.

The effect sizes reported here are surrounded by statistical uncertainty, and our back-of-the-envelope calculations omit global general equilibrium effects operating through international prices or trade. Our calculations also omit other benefits of the Green Revolution, such as improvements in health associated with greater food availability; we also omit potential costs, such as environmental damage. Still, the numbers strongly suggest that the development and diffusion of HYVs has been an important source of economic growth in developing countries.

To put our estimated effect sizes into perspective, the effect of delaying the Green Revolution by 10 years is of a magnitude comparable (with opposite sign) to the income effect of democratizing, which Acemoglu et al. (2019) estimate to be about 20% after 25 years, and to the effect of railroad access in nineteenth-century India, which Donaldson (2018) puts at 16%. The population effect we find is substantially smaller than the effect of medical innovations, which, according to Acemoglu, Fergusson, and Johnson (2020), has increased the population by 45% between 1940 and 1980 in their sample of countries and by even more in low- and middle-income countries.

Considerable heterogeneity is hidden beneath the aggregate effects of the Green Revolution discussed above, as country-level impacts depended on agroecology and the initial size of the agricultural sector. The heterogeneity is visible in figure 3, but in table 6 we quantify it for selected subsamples, using our baseline country-level estimates. We should emphasize that the calculations disregard possible treatment heterogeneity across subgroups, meaning that the heterogeneity we uncover in this exercise reflects heterogeneity in the exposure to the Green Revolution alone. The results show that the third of countries in our sample with the highest exposure to the Green Revolution would have been 58% poorer in the counterfactual without the Green Revolution, whereas the least affected third would be "only" 31% poorer. We do not find large differences across income groups, as measured by GDP per capita in 1964.

The Green Revolution is often associated with the 1960s and 1970s, but rather than slowing down, the rate of adoption and the number of new HYVs increased in the 1980s, 1990s, and 2000s. Scattered evidence from sub-Saharan Africa suggests that the HYV adoption rate has increased by as much in the 2000s as in the four preceding decades.[27] One reason is that, compared to that in other parts of the world, especially Southeast Asia, African agriculture is specialized in cassava, sorghum, millet, and other crops for which HYVs became available relatively late. Our results consequently shed light on the divergence between Southeast Asia and Africa during the second half of the twentieth century.

The growth effect of increasing agricultural productivity naturally declines with the size of the agricultural sector relative to GDP. The contribution to aggregate income growth from further investments in agricultural research will therefore be smaller in the future, as agriculture shrinks as a share of the global economy. Yet agriculture still accounts for about 40% of employment in the average developing country, and the technological frontier continues to shift outward—not only for yield increases but also for environmental benefits and resilience to climate change. Our results suggest that investments in the development and diffusion of agricultural technology have substantially improved living standards in the poorest places on our planet over the past half-century. Further investments in agricultural science targeting the developing world may have the potential to sustain these gains in the decades ahead.

References

Acemoglu, Daron, and Joshua Angrist. 2000. "How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws." *NBER Macroeconomics Ann.* 15:9–59.

Acemoglu, Daron, Leopoldo Fergusson, and Simon Johnson. 2020. "Population and Conflict." *Rev. Econ. Studies* 87 (4): 1565–604.

Acemoglu, Daron, and Simon Johnson. 2007. "Disease and Development: The Effect of Life Expectancy on Economic Growth." *J.P.E.* 115 (6): 925–85.

Acemoglu, Daron, Suresh Naidu, Pascual Restrepo, and James A. Robinson. 2019. "Democracy Does Cause Growth." *J.P.E.* 127 (1): 47–100.

Alsan, Marcella, and Claudia Goldin. 2019. "Watersheds in Child Mortality: The Role of Effective Water and Sewerage Infrastructure, 1880–1920." *J.P.E.* 127 (2): 586–638.

Angelsen, Arild and David Kaimowitz. 2001. "Introduction: The Role of Agricultural Technologies in Tropical Deforestation." In *Agricultural Technologies and Tropical Deforestation*, edited by Arild Angelsen and David Kaimowitz, 1–17. Wallingford: CABI.

[27] Calculations are based on data from the CGIAR's Diffusion and Impact of Improved Varieties in Africa (DIIVA) project.

Autor, David H. 2003. "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing." *J. Labor Econ.* 21 (1): 1–42.

Barker, Randolph, Robert W. Herdt, and Beth Rose. 1985. *The Rice Economy of Asia*, volume 2. Washington, DC: Resources for the Future and Manila: Internat. Rice Res. Inst.

Bartik, Alexander W., Janet Currie, Michael Greenstone, and Christopher R. Knittel. 2019. "The Local Economic and Welfare Consequences of Hydraulic Fracturing." *American Econ. J.: Appl. Econ.* 11 (4): 105–55.

Becker, Gary S., Kevin M. Murphy, and Robert Tamura. 1990. "Human Capital, Fertility, and Economic Growth." *J.P.E.* 98 (5, pt. 2): S12–S37.

Bharadwaj, Prashant, James Fenske, Namrata Kala, and Rinchan Ali Mirza. 2020. "The Green Revolution and Infant Mortality in India." *J. Health Econ.* 71:102314.

Bhuller, Manudeep, Tarjei Havnes, Edwin Leuven, and Magne Mogstad. 2013. "Broadband Internet: An Information Superhighway to Sex Crime?" *Rev. Econ. Studies* 80 (4): 1237–66.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global Non-linear Effect of Temperature on Economic Production." *Nature* 527 (7577): 235.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli. 2016. "Agricultural Productivity and Structural Transformation: Evidence from Brazil." *A.E.R.* 106 (6): 1320–65.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Rev. Econ. and Statis.* 90 (3): 414–27.

Card, David. 2001. "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *J. Labor Econ.* 19 (1): 22–64.

Córdoba, Juan Carlos, and Marla Ripoll. 2009. "Agriculture and Aggregation." *Econ. Letters* 105 (1): 110–12.

Costinot, Arnaud, and Dave Donaldson. 2011. "How Large Are the Gains from Economic Integration? Theory and Evidence from US Agriculture, 1880–2002." Manuscript, Massachusetts Inst. Technology.

Dalrymple, Dana G. 1974. "Development and Spread of High-Yielding Varieties of Wheat and Rice in the Less Developed Nations." Foreign Agricultural Econ. Report no. 95, Econ. Res. Service, US Dept. Agriculture, Washington, DC.

———. 1985. "The Development and Adoption of High-Yielding Varieties of Wheat and Rice in Developing Countries." *American J. Agricultural Econ.* 67 (5): 1067–73.

———. 1986. *Development and Spread of High-Yielding Rice Varieties in Developing Countries.* Washington, DC: Agency for Internat. Development.

———. 2008. "International Agricultural Research as a Global Public Good: Concepts, the CGIAR Experience and Policy Issues." *J. Internat. Development* 20 (3): 347–79.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Econ. J.: Macroeconomics* 4 (3): 66–95.

Donaldson, Dave. 2018. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *A.E.R.* 108 (4–5): 899–934.

Duflo, Esther, Michael Greenstone, Raymond Guiteras, and Thomas Clasen. 2015. "Toilets Can Work: Short and Medium Run Health Impacts of Addressing Complementarities and Externalities in Water and Sanitation." Working Paper no. 20521 (September), NBER, Cambridge, MA.

Evenson, Robert E., and Douglas Gollin. 2003a. "Assessing the Impact of the Green Revolution, 1960 to 2000." *Science* 300 (5620): 758–62.

———, eds. 2003b. *Crop Variety Improvement and Its Effect on Productivity: The Impact of International Agricultural Research.* Wallingford: CABI.

Evenson, Robert E., and Mark W. Rosegrant. 2003. "The Economic Consequences of Crop Genetic Improvement Programmes." In *Crop Variety Improvement and Its Effect on Productivity: The Impact of International Agricultural Research,* edited by Robert E. Evenson and Douglas Gollin, 473–97. Wallingford: CABI.

Foster, Andrew D., and Mark R. Rosenzweig. 2004. "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000." *Econ. Development and Cultural Change* 52 (3): 509–42.

———. 2007. "Economic Development and the Decline of Agricultural Employment." In *Handbook of Development Economics,* vol. 4, edited by T. Paul Schultz and John A. Strauss, 3051–83. Amsterdam: North-Holland.

Galor, Oded, and Ömer Özak. 2016. "The Agricultural Origins of Time Preference." *A.E.R.* 106 (10): 3064–103.

Galor, Oded, and David N. Weil. 2000. "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond." *A.E.R.* 90 (4): 806–28.

Godfray, H. Charles J., and Sherman Robinson. 2015. "Contrasting Approaches to Projecting Long-Run Global Food Security." *Oxford Rev. Econ. Policy* 31 (1): 26–44.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *A.E.R.* 110 (8): 2586–624.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "The Agricultural Productivity Gap." *Q.J.E.* 129 (2): 939–93.

Gollin, Douglas, Stephen Parente, and Richard Rogerson. 2002. "The Role of Agriculture in Development." *A.E.R.* 92 (2): 160–64.

———. 2007. "The Food Problem and the Evolution of International Income Levels." *J. Monetary Econ.* 54 (4): 1230–55.

Goodman-Bacon, Andrew. 2018. "Difference-in-Differences with Variation in Treatment Timing." Working Paper no. 25018 (September), NBER. Cambridge, MA.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen. 2020. "Do Credit Market Shocks Affect the Real Economy? Quasi-Experimental Evidence from the Great Recession and 'Normal' Economic Times." *American Econ. J.: Econ. Policy* 12 (1): 200–225.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (4): 501–22.

Hornbeck, Richard, and Pinar Keskin. 2014. "The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought." *American Econ. J.: Appl. Econ.* 6 (1): 190–219.

Houthakker, H. S. 1955. "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis." *Rev. Econ. Studies* 23 (1): 27–31.

Lewis, W. Arthur. 1951. *Principles of Economic Planning.* London: Dennis Dobson.

Maredia, Mywish K., and Derek R. Byerlee. 2000. "Efficiency of Research Investments in the Presence of International Spillovers: Wheat Research in Developing Countries." *Agricultural Econ.* 22 (1): 1–16.

Matsuyama, Kiminori. 1992. "Agricultural Productivity, Comparative Advantage, and Economic Growth." *J. Econ. Theory* 58 (2): 317–34.

Mitchell, Brian R. 1982. *International Historical Statistics: Africa and Asia.* New York: New York Univ. Press.

Moscona, Jacob. 2019. "Agricultural Development and Structural Change, within and across Countries." Manuscript.

Nunn, Nathan, and Nancy Qian. 2011. "The Potato's Contribution to Population and Urbanization: Evidence from a Historical Experiment." *Q.J.E.* 126 (2): 593–650.

Nurkse, Ragnar. 1953. *Problems of Capital Formation in Underdeveloped Countries.* Oxford: Oxford Univ. Press.

Pandey, Sushil, Maria Lourdes Velasco, and Takashi Yamano. 2015. "Scientific Strength in Rice Improvement Programmes, Varietal Outputs and Adoption of Improved Varieties in South Asia." In *Crop Improvement, Adoption, and Impact of Improved Varieties in Food Crops in Sub-Saharan Africa,* edited by Thomas S. Walker and Jeffrey Alwang, 239–64. Wallingford: CABI.

Perez, Nicostrato D., and Mark W. Rosegrant. 2015. "The Impact of Investment in Agricultural Research and Development and Agricultural Productivity." IFPRI Discussion Paper no. 1447, Internat. Food Policy Res. Inst., Washington, DC.

Perkins, John H. 1997. *Geopolitics and the Green Revolution: Wheat, Genes, and the Cold War.* New York: Oxford Univ. Press.

Pingali, Prabhu, and Tim Kelley. 2007. "The Role of International Agricultural Research in Contributing to Global Food Security and Poverty Alleviation: The Case of the CGIAR." In *Handbook of Agricultural Economics,* vol. 3, edited by Robert E. Evenson and Prabhu Pingali, 2381–418. Amsterdam: North-Holland.

Pingali, Prabhu L. 2012. "Green Revolution: Impacts, Limits, and the Path Ahead." *Proc. Nat. Acad. Sci. USA* 109 (31): 12302–08.

Raitzer, David A., and Timothy G. Kelley. 2008. "Benefit-Cost Meta-analysis of Investment in the International Agricultural Research Centers of the CGIAR." *Agricultural Systems* 96 (1–3): 108–23.

Renkow, Mitch, and Derek Byerlee. 2010. "The Impacts of CGIAR Research: A Review of Recent Evidence." *Food Policy* 35 (5): 391–402.

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu. 2008. "Agriculture and Aggregate Productivity." *J. Monetary Econ.* 55 (2): 234–50.

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb. 2019. "Fast and Wild: Bootstrap Inference in Stata Using Boottest." *Stata J.* 19 (1): 4–60.

Rose, Beth. 1985. *Appendix to The Rice Economy of Asia.* Washington, DC: Resources for the Future and Manila: Internat. Rice Res. Inst.

Rosenstein-Rodan, P. N. 1943. "Problems of Industrialisation of Eastern and South-Eastern Europe." *Econ. J.* 53 (210–211): 202–11.

Rusike, J., N. Mahungu, S. Jumbo, V. Sandifolo, and G. Malindi. 2010. "Estimating Impact of Cassava Research for Development Approach on Productivity, Uptake and Food Security in Malawi." *Food Policy* 35 (2): 98–111.

Thirtle, Colin, Lin Lin, and Jenifer Piesse. 2003. "The Impact of Research-Led Agricultural Productivity Growth on Poverty Reduction in Africa, Asia and Latin America." *World Development* 31 (12): 1959–75.

Vollrath, Dietrich. 2011. "The Agricultural Basis of Comparative Development." *J. Econ. Growth* 16 (4): 343–70.

von der Goltz, Jan, Aaditya Dar, Ram Fishman, Nathaniel D. Mueller, Prabhat Barnwal, and Gordon C. McCord. 2020. "Health Impacts of the Green Revolution: Evidence from 600,000 Births across the Developing World." *J. Health Econ.* 74:102373.

Walker, Thomas S., and Jeffrey Alwang, eds. 2015. *Crop Improvement, Adoption and Impact of Improved Varieties in Food Crops in Sub-Saharan Africa.* Wallingford: CABI.