Two Blades of Grass: The Impact of the Green Revolution*

Douglas Gollin Casper Worm Hansen Asger Wingender†

February 4, 2021

Abstract

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 percent between 1965 and 2010 with further gains coming through reallocation of inputs. Higher yields increased income and reduced population growth. A ten-year delay of the Green Revolution would in 2010 have cost 17 percent of GDP/capita and added 223 million people to the developing world population. The cumulative GDP loss would have been US$83 trillion, corresponding to one year of current global GDP.

Keywords: Green Revolution; High Yielding Variety Crops; Productivity Shock; Macroeconomic Development. **JEL:** N50; O11; O13; O50; Q16.

*This paper is a substantially revised version of earlier working papers with the same title but different subtitles. Acknowledgments: We would like to thank the editor (Michael Greenstone), anonymous referees, Philipp Ager, Samuel 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 for useful comments and suggestions from seminar participants at the University of Copenhagen, McGill University, Paris School of Economics, the IMF Research Department, the 7th Annual Workshop on Growth, History and Development at the University of Southern Denmark, the 2016 annual EEA meeting in Geneva, the 2017 annual AEA meeting in Chicago, the NBER conference on Trade and Agriculture, the CEPR conference: Accounting For The Wealth Of Nations: History And Theory, and the Brown University conference on Deep Determinants of Economic Development. Early drafts of this paper were written while Douglas Gollin was serving on the Standing Panel on Impact Assessment (SPIA) of the CGIAR, the publicly funded agricultural research and development consortium under whose auspices many of the Green Revolution varieties were developed. Gollin's views do not represent those of SPIA or the CGIAR.

†Contact: Douglas Gollin, Oxford University, douglas.gollin@qeh.ox.ac.uk; Casper Worm Hansen, University of Copenhagen, Casper.Worm.Hansen@econ.ku.dk; Asger Wingender, University of Copenhagen, amw@econ.ku.dk

> "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 in Gulliver's Travels*

1 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; or 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 impacted 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

[1] See, for example, Córdoba and Ripoll (2009); Gollin et al. (2002, 2007); Restuccia et al. (2008); Vollrath (2011).

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 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 will be 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., Bustos et al. 2016 and Hornbeck and Keskin 2014) 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 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 pre-existing 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 impacted 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 percent between 1965 and 2010. The total effect on yields is even higher because of substitution towards 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 per capita. Our estimates imply that delaying the Green Revolution for ten years would have reduced GDP per capita in 2010 by US$1,273 (PPP adjusted), or 17 percent, 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 ten 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 percent 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 20th 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 ten 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

[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 et al. (2001, p.3)

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 crops (Bustos et al., 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 von der Goltz et al. (2020), Bharadwaj et al. (2020), and Moscona (2019).

2 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 will 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 been available to breeders. Our second claim is that the *timing* of the initial research was

[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 et al. (2003); Pingali and Kelley (2007); Dalrymple (2008); Raitzer and Kelley (2008); Rusike et al. (2010); Costinot and Donaldson (2011).

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.

2.1 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 non-food 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]

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 Baï¿œos 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 previously in a number of sources, e.g., Dalrymple et al. (1974); Dalrymple (1985, 1986); Barker et al. (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

[4]It is true that plant breeding took place prior to the Green Revolution within many national programs. But prior to 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 years of the International Rice Research Institute, breeders averaged over 2500 crosses per year.

[5]IRRI too had a modest precursor program, a small breeding effort initiated under the auspices of the 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.

support investments in rice research; nor that it would support a wheat research program that was based in Mexico.

Against this backdrop, rice breeding began at IRRI in 1965. Within the first weeks of breeding effort, scientists made a cross (designated IR8) that gave rise to what would eventually prove to be the first "mega-variety" of rice. The other research center, CIMMYT, which was built on an earlier local research program in Mexico, 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 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

[6]There are now fifteen 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, non-governmental organizations, private philanthropies, and other donors.

South Asia (Pandey et al., 2015).

2.2 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 some 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 by the authors based on historical records of varietal releases and on other analysis of breeding data, shows the release dates of HYVs for different crops; detailed documentation is in Online Appendix Table A1.

Table 1 about here

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 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 which significantly reduces yields. The first rust-resistant variety was released in 1948, and further improved varieties followed in the 1950s. Work on semi-dwarfism began in the mid 1950s, and culminated with the release of the first truly high-yielding varieties in Mexico.[7] Based on the success of these varieties, Norman

[7]Semi-dwarf varieties of rice and wheat were central to the Green Revolution. Because they were shorter and stiffer-strawed, 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.

Borlaug, the lead scientist at the Mexican research center brought seeds of the initial semi-dwarf wheats to India in 1963. In 1965, two semi-dwarf varieties originating from the Mexican program were released in India; more or less concurrently, semi-dwarf 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 to 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 semi-dwarf 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 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.

Figure 1 about here

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.

3 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 $\bar{X}_k$, $k = 1, 2, .., N$. Regions are not necessarily geographically distinct, but 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 only grow a single crop k. Land within each region is divided into infinitesimally small plots, indexed by $i \in \left[0, \bar{X}_k\right]$ 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 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,

[8] We only report the key equations, but the intermediate steps are provided in Online Appendix B.

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 all plots $i > X_{kt}$ are left fallow. Moreover, let a_{ik} distributed smoothly and differentiably across plots in such a way that the average productivity of crop land within a region is given by:

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

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]

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

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}:

$$yield_{kt} = \frac{Y_{Akt}}{X_{kt}} = a_{kt} s_{kt}^{-\delta}. \tag{3}$$

To highlight the effect of HYVs, assume that the crop-specific (and region-specific) productivity

[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 Online Appendix Equation A2 for the intermediate steps needed to derive this equation.

parameter a_{kt} can be decomposed as $a_{kt} = a_{kt}^{HYV} \bar{a}_k z_t u_{kt}$, where $a_{kt}^{HYV} = 1$ if no HYVs are available, and $a_{kt}^{HYV} > 1$ 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:

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

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

We now turn to the aggregate implications of introducing HYVs. To simplify notation, we abstract from other factors affecting yields other than HYVs 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:

$$GR_t = \sum_{k=1}^{N} a_{kt}^{HYV} \frac{Y_{Ak0}}{Y_{A0}}, \tag{5}$$

where $\frac{Y_{Ak0}}{Y_{A0}}$ is the share of crop k in aggregate crop production prior to the Green Revolution. In our empirical analysis, we use Equation 5 and our estimates of $\ln a_{kt}^{HYV}$ to construct an empirical counterpart of GR_t, which allows us to identify country-level effects of the Green Revolution. A value of GR_t of 1.5 implies that HYVs, everything else being equal, have increased aggregate yields by 50 percent. 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:

$$\Delta \ln yield_t = \ln GR_t + \ln \Phi_t - \delta \Delta \ln X_t, \tag{6}$$

where $\Phi_t \equiv \frac{\left(\sum_{k=1}^{N}\left(a_{kt}^{HYV}\bar{a}_k\right)^{\frac{1}{\delta}}\bar{X}_k\right)^{\delta}}{\sum_{k=1}^{N} a_{kt}^{HYV}\bar{a}_k^{\frac{1}{\delta}}\bar{X}_k}$. 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 ($\ln GR_t$), a productivity gain from reallocation of crop land toward regions growing HYVs ($\ln \Phi_t$), and changes to the extent of crop land ($\delta \Delta \ln X_t$), which affect yields because of decreasing returns to scale. The two first effects are unambiguously positive, whereas the latter 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 $\bar{c}_a$, 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 $yield_t \cdot X_t$. This equilibrium condition allows us to write yield growth as:

$$\Delta \ln yield_t = \frac{1}{1-\delta}\ln GR_t + \frac{1}{1-\delta}\ln \Phi_t - \frac{\delta}{1-\delta}\Delta \ln L_t, \tag{7}$$

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 et al. (1990) and Galor and Weil (2000).

Our discussion has so far only dealt 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:

$$y_t = \frac{Y_t}{L_t} = \frac{Y_A}{L} + \frac{Y_M}{L} = \bar{c}_a + p_t m \frac{L_{Mt}}{L_t}\,, \tag{8}$$

where p_t is the relative price of manufactured goods. This equation shows that structural transformation is instrumental in order to increase GDP per capita, and our assumptions mean that structural transformation ultimately is driven by yields (because $\Delta \ln \frac{L_{At}}{L_t} = -\Delta \ln yield_t$). 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:

$$\Delta \ln y_t \approx \frac{p_0 m}{y_0} \frac{L_{A0}}{L_0} \Delta \ln yield_t. \tag{9}$$

A shock to yields is not only moderated by the size of the labor force initially employed in agriculture, but also by the term $\frac{p_0 m}{y_0}$, 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 et al., 2014), $\frac{p_0 m}{y_0} > 1$, 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.

4 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 roll-out) 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_t 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.

4.1 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:

$$\ln yield_{kit} = \sum_k \sum_{j \in T_k} \beta_{kj} \cdot \mathbf{1}_{kt}^{t=\tau_k+j} + \delta \ln harea_{kit} + \mu_{ik} + \mu_{it} + \varepsilon_{kit}, \tag{10}$$

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

[10] Staggered adoption (roll-out) 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 et al., 2020).

effects control 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 2, 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 $\ln a_{kt}^{HYV}$ in the theoretical framework outlined in Section 3 (see also Online Appendix 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, $\ln harea_{kit}$, 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.,

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

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:

$$\ln yield_{kit} = \sum_{k} \alpha_k \cdot \mathbf{1}_{kt}^{HYV} \cdot t + \delta \ln harea_{kit} + \mu_{ki} + \mu_{it} + \varepsilon_{kit}, \tag{11}$$

where $\mathbf{1}_{kt}^{HYV}$ is an indicator equal to one in years after the release of the first HYV of a crop k. For pure control crops, $\mathbf{1}_{kt}^{HYV}$ 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, i.e, $\alpha_k = \alpha$, and common trends within crop groups, i.e., $\alpha_k = \alpha_\kappa$ 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.

[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 Panel F of Online Appendix Figure A6.

4.2 Country-level framework

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

$$\widehat{GR}_{it} = \sum_{k=1}^{N} \exp\left(\hat{\alpha}_k \mathbf{1}_{kt}^{HYV} t\right) \frac{Y_{Aik0}}{Y_{Ai0}}, \tag{12}$$

where the observed pre-Green Revolution production shares, Y_{Ak0}/Y_{A0}, are measured in constant prices and averaged over 1961-1964 to reduce noise from, e.g, weather shocks. As above, the indicator $\mathbf{1}_{kt}^{HYV}$ is zero the entire period for crops of which HYVs have not been developed. The log of $\widehat{GR}_{it}$ is analogous to a shift-share instrument for log yields. Because $\ln \widehat{GR}_{it}$ by construction is zero before the onset of the Green Revolution, we can interpret $\ln \widehat{GR}_{it}$ 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 $\ln \widehat{GR}_{it}$ as predicted GR yields. To 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:

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

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 3, 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 Equation 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 et al. (2008), adapted to our setting.[13]

[13] Our implementation of the bootstrap resembles the procedure for 2SLS regressions with clustered standard errors described by Roodman et al. (2019). What is different in our setting is that our baseline

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:

$$\ln y_{it} = \lambda(\ln \widehat{GR}_{it}) \times \frac{L_{Ai0}}{L_{i0}} + \mathbf{X}_{it}\rho + \mu_i + \mu_t + \epsilon_{it} \tag{14}$$

where L_{Ai0}/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 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 et al. (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

regressions are reduced form rather than 2SLS, and that our reduced form generated regressor is aggregated from first step estimates using Equation 12. We report bootstrapped standard errors based on 1000 replications. Both Cameron et al. (2008) and Roodman et al. (2019) recommend inference based on bootstrapped p-values, which in some cases have 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.

that countries were producing before the Green Revolution. Another way of specifying this variation is, as Goldsmith-Pinkham et al. (2018) argue in their discussion of shift-share instruments, to use the initial production shares Y_{Aik0}/Y_{Ai0} as separate instruments. The relevance of each instrument is essentially our estimated effect of HYVs on crop yields, i.e., $\hat{\alpha}_k$, making our identification strategy equivalent to a difference-in-difference strategy with a continuous treatment intensity measured by Y_{Aik0}/Y_{Ai0}. 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 pre-trends 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 only started to emerge 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:

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

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.

5 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 the League of Nations (various issues), the Statistical Yearbook of the United Nations (various issues), reports from the Economic Research Service of the U.S. Department of Agriculture, and from 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 1920-1961. 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

[14]Ideally, we would use country-level prices for all crops prior to the Green Revolution. Comprehensive farm-gate price data in local currency are available from FAO, but the data start in 1966, only cover 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 we use in the aggregation (see Online Appendix Section 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 Online Appendix Table A2. See Online Appendix Figure A5 for a comparison with prices.

[15]There are 19 IARC mandate crops, but we exclude plantains, potatoes, and soybeans from the sample

and to a set of botanically similar pure control crops for which no breeding of HYVs have taken place (see Section 4.1). 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 percent 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.

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 cut-off 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 the private sector, which markets genetically modified (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.

[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.

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 with fewer than 10,000 hectares of arable land devoted to food production. We additionally exclude the ten 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]

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 Online Appendix D.

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

6 Results

6.1 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 non-parametric trends in relative yields of a crop before and after the first HYV of the crop was released. Panel A displays estimates under the assumption that HYVs have identical effects on yields of all 16 crops of which HYVs were released. For event years before the first HYV release (the period up until the vertical red line), the estimated coefficients are close to zero and statistically insignificant, implying that yields of treated crops and untreated crops followed the same trend prior to the first HYV release. The absence of differential pre-trends 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.

Figure 2 about here

Converting our estimates from logarithms to levels, we find that relative yields are on average nine percent higher 10 years after a HYV release ($\beta_{10} = 0.09$), and 75 percent 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 percent 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.

Panels B-D of Figure 2 report separate event studies for cereals, pulses, and roots and tubers, thereby allowing for yield gains to differ for botanically different crops. For all three crop types, we find no evidence of pre-trends, and highly significant increases in relative yields following the 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 non-parametric 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 percent larger than the gain for roots and tubers, and 85 percent larger than the gain for pulses. At the level of individual crops (see Online Appendix 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 following the Green Revolution. Yield gains for sorghum and millet have, for instance, been modest.

Table 2 about here

Columns 3-7 show the robustness of the results in Column 2. In Column 3, we exclude observations with low data quality.[20] The estimates are slightly higher than our baseline, suggesting attenuation bias is present, but small. In Column 4, we disregard the historical data we have collected and only use 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 pre-trends from the estimates. Unsurprisingly, given the

[19]Online Appendix Figures A1 and A2 report separate event studies for each of the treated crops in the categories cereals, and roots and 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, and official data that are 1) unvarying from year to year, 2) have less than three significant digits (we remove all zeroes trailing the last non-zero digit and count the number of digits left), and 3) are based on crops for which harvested area is less than 1000 hectares. The two first restrictions eliminate data points that are crude estimates, the last restriction removes outliers from small sample sizes in the underlying surveys.

picture in Figure 2, we find estimates highly similar to those in Column 2.

In Online Appendix Table A6, we demonstrate that the harvested area of treated crops increases relative to that of untreated crops, meaning that farmers substitute towards 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 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.

6.2 The spatial distribution of the Green Revolution

Armed with our crop-specific difference-in-difference estimates, we use Equation 12 to construct $\widehat{GR}_{it}$, i.e., 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 percent 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 Panel A of Figure 3 shows the most important food crop in each country in our sample, measured as the value of production in 1961 (in 1966 prices). Unsurprisingly, South-East Asia shows up in this figure as rice territory, whereas there is more within-region variation in for 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 $\ln \widehat{GR}_{it}$ in a map as we do in Panel B

Figure 3 about here

of Figure 3. 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 Online Appendix Figure A3 for an illustration of the time variation). Darker hues of green in the map indicate 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 South-East Asia, despite rice being the most important crop in all but a few countries in the region. Still, there is less variation within South-East Asia than within Africa, where there is more heterogeneity in the crop mix.

6.3 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 2-6. The robustness checks show that we for all five outcomes obtain results similar to the baseline if we aggregate yields using caloric content rather than prices (Column 2), control for weather shocks and climate change (Column 3), control for initial yield levels interacted with year fixed effects (Column 4), and exclude IARC host countries (Column 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, amplify the *ceteris paribus* effect of HYVs by 73 percent. The estimated elasticity is only significantly larger than unity at the 10 percent 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 crop land, reported in Panel C, is smaller in magnitude, as land devoted to non-food crops (e.g., cotton, tobacco) was not directly affected by the introduction of HYVs of food crops.

Table 3 about here

In Panel D, we look at the adoption of HYVs, measured by the share of crop land devoted to them. The data, taken from Evenson and Gollin (2003a), end in 2000 and only cover 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 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, which, 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]

6.4 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

[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 Online Appendix Figure A6.

simultaneously. HYVs of other crops only began to emerge more than a decade later, and with a few exceptions, the yield gains for these late-comers were generally smaller than for the 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 pre-trends that might invalidate our research design.

Figure 4 about here

The event-study for GDP per capita in Panel A of Figure 4 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 pre-trend in the point estimates, but as shown in Panel B, controlling for pre-Green Revolution GDP growth interacted with year fixed effects eliminates the pre-trend while leaving the post-Green Revolution estimates intact. The event-study graph for population growth, in Panel C of Figure 4, 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, Panel 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 pre-trends, 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 at the demographic transition before the Green Revolution. As a result, the pre-trend is eliminated, and the effect of HYVs on population size becomes slightly stronger.[23]

[23] Online Appendix Figure 7A 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 Figure 4. Online Appendix Figures 8A 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

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 $\ln \widehat{GR}_{it}$ 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, Column 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 in Table 3 found 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 towards 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 in Panel C find 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 Panel 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.

negative effect on population is most pronounced for rice-growing countries. There are no visible pre-trends for any of the three crops.

[24]Event-study graphs corresponding to this specification are reported in Online Appendix Figure 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 (Online Appendix Figure A11). This finding, which is consistent with the model in Section 3, shows that it is unnecessary to interact with the size of the agricultural sector when studying agricultural outcomes.

Table 4 about here

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 pre-trends. The other robustness checks mimic those in Table 3. The results show that our baseline results are robust to these alternative specifications, including to controlling for climate change, which has been found to have a negative impact on economies of the developing world (e.g., Dell et al., 2012 and Burke et al., 2015), and obviously also affect agriculture. We do, however, find slightly larger effects on population size when pre-trends are controlled for. Further robustness checks are reported in Online Appendix H.

The demographic response to the Green Revolution is clearly of first-order importance in order to explain 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 one percent 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 Column 2). Lower mortality would, by itself, increase the population size, so the impact on fertility (or migration) needs to 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 on population is entirely driven by changes in mortality and fertility. The last column shows

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

positive effects on human capital, consistent with a quantity-quality trade-off. In Appendix Table A13, we extend the analysis to a range of additional indicators of development.

Table 5 about here

7 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, respectively. 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]

Table 6 about here

[26] A detailed description of these calculations can be found in Online Appendix Section I.

Our baseline estimates imply that aggregate food crop yields for our sample of countries would have been 49 percent lower in 2010 had the Green Revolution never happened. They would still have been higher than 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 percent 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 have been 2.3 percent 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 ten-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 percent 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 four percent higher, corresponding to 223 million people, had the Green Revolution been delayed by ten 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 ten years is of a comparable magnitude (with opposite sign) to the income effect of democratizing, which Acemoglu et al. (2019) estimate to be about 20 percent after 25 years, and to the effect of railroad access in 19th century India, which Donaldson (2018) puts at 16 percent. The population effect we find is substantially smaller than the effect of medical innovations, which according to Acemoglu et al. (2020) has increased the population by 45 percent 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 agro-ecology and the initial size of the agricultural sector. The heterogeneity is visible in Figure 3 above, but in Table 6 we quantify it for selected sub-samples using our baseline country-level estimates. We should emphasize that the calculations disregard possible treatment heterogeneity across sub-groups, 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 percent poorer in the counterfactual without the Green Revolution, whereas the least affected third would "only" be 31 percent 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 other parts of the world, especially South-East 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 South-East Asia and Africa during the second half of the 20th 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 percent 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, D., Angrist, J., 2000. How large are human-capital externalities? evidence from compulsory schooling laws. NBER macroeconomics annual 15, 9–59.

Acemoglu, D., Fergusson, L., Johnson, S., 2020. Population and conflict. The Review of Economic Studies 87, 1565–1604.

Acemoglu, D., Johnson, S., 2007. Disease and development: the effect of life expectancy on economic growth. Journal of Political Economy 115, 925–985.

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

Acemoglu, D., Naidu, S., Restrepo, P., Robinson, J.A., 2019. Democracy does cause growth. Journal of Political Economy 127, 47–100.

Alsan, M., Goldin, C., 2019. Watersheds in child mortality: The role of effective water and sewerage infrastructure, 1880–1920. Journal of Political Economy 127.

Angelsen, A., Kaimowitz, D., et al., 2001. Agricultural technology and forests: a recapitulation. Agricultural technologies and tropical deforestation .

Autor, D.H., 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. Journal of labor economics 21, 1–42.

Barker, R., Herdt, R.W., Rose, B., 1985. The rice economy of Asia. volume 2. Int. Rice Res. Inst.

Bartik, A.W., Currie, J., Greenstone, M., Knittel, C.R., 2019. The local economic and welfare consequences of hydraulic fracturing. American Economic Journal: Applied Economics 11, 105–55.

Becker, G.S., Murphy, K.M., Tamura, R., 1990. Human capital, fertility, and economic growth. Journal of political economy 98, S12–S37.

Bharadwaj, P., Fenske, J., Kala, N., Mirza, R.A., 2020. The green revolution and infant mortality in india. Journal of Health Economics 71, 102314.

Bhuller, M., Havnes, T., Leuven, E., Mogstad, M., 2013. Broadband internet: An information superhighway to sex crime? Review of Economic Studies 80, 1237–1266.

Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. Nature 527, 235.

Bustos, P., Caprettini, B., Ponticelli, J., 2016. Agricultural productivity and structural transformation: Evidence from brazil. American Economic Review 106, 1320–65.

Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. The Review of Economics and Statistics 90, 414–427.

Card, D., 2001. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. Journal of Labor Economics 19, 22–64.

Córdoba, J.C., Ripoll, M., 2009. Agriculture and aggregation. Economics Letters 105, 110–112.

Costinot, A., Donaldson, D., 2011. How large are the gains from economic integration? theory and evidence from us agriculture, 1880-2002. Massachusetts Institute of Technology. Mimeo .

Dalrymple, D.G., 1985. The development and adoption of high-yielding varieties of wheat and rice in developing countries. American Journal of Agricultural Economics 67, 1067–1073.

Dalrymple, D.G., 1986. Development and spread of high-yielding rice varieties in developing countries. Int. Rice Res. Inst.

Dalrymple, D.G., 2008. International agricultural research as a global public good: concepts, the CGIAR experience and policy issues. Journal of International Development 20, 347–379.

Dalrymple, D.G., et al., 1974. Development and spread of high-yielding varieties of wheat and rice in the less developed nations. Technical Report. United States Department of Agriculture, Economic Research Service.

Dell, M., Jones, B.F., Olken, B.A., 2012. Temperature shocks and economic growth: Evidence from the last half century. American Economic Journal: Macroeconomics 4, 66–95.

Donaldson, D., 2018. Railroads of the raj: Estimating the impact of transportation infrastructure. American Economic Review 108, 899–934.

Duflo, E., Greenstone, M., Guiteras, R., Clasen, T., 2015. Toilets can work: Short and medium run health impacts of addressing complementarities and externalities in water and sanitation. Technical Report. National Bureau of Economic Research.

Evenson, R.E., Gollin, D., 2003a. Assessing the impact of the green revolution, 1960 to 2000. science 300, 758–762.

Evenson, R.E., Gollin, D., 2003b. Crop variety improvement and its effect on productivity:

The impact of international agricultural research. Cabi.

Evenson, R.E., Rosegrant, M., 2003. The economic consequences of crop genetic improvement programmes. Crop variety improvement and its effect on productivity: the impact of international agricultural research , 473–497.

Foster, A.D., Rosenzweig, M.R., 2004. Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970–2000. Economic Development and Cultural Change 52, 509–542.

Foster, A.D., Rosenzweig, M.R., 2007. Chapter 47 economic development and the decline of agricultural employment, Elsevier. volume 4 of *Handbook of Development Economics*, pp. 3051 – 3083.

Galor, O., Özak, Ö., 2016. The agricultural origins of time preference. American Economic Review 106, 3064–3103.

Galor, O., Weil, D.N., 2000. Population, technology, and growth: From malthusian stagnation to the demographic transition and beyond. American economic review 90, 806–828.

Godfray, H.C.J., Robinson, S., 2015. Contrasting approaches to projecting long-run global food security. Oxford Review of Economic Policy 31, 26–44.

Goldsmith-Pinkham, P., Sorkin, I., Swift, H., 2018. Bartik instruments: What, when, why, and how. Technical Report. National Bureau of Economic Research.

Gollin, D., Lagakos, D., Waugh, M.E., 2014. The agricultural productivity gap. The Quarterly Journal of Economics 129, 939–993.

Gollin, D., Parente, S.L., Rogerson, R., 2002. The role of agriculture in development. American Economic Review Papers and Proceedings 92.

Gollin, D., Parente, S.L., Rogerson, R., 2007. The food problem and the evolution of international income levels. Journal of Monetary Economics 54, 1230–1255.

Goodman-Bacon, A., 2018. Difference-in-differences with variation in treatment timing. Technical Report. National Bureau of Economic Research.

Greenstone, M., Mas, A., Nguyen, H.L., 2020. Do credit market shocks affect the real economy? quasi-experimental evidence from the great recession and" normal" economic times. American Economic Journal: Economic Policy 12, 200–225.

Griliches, Z., 1957. Hybrid corn: An exploration in the economics of technological change. Econometrica, Journal of the Econometric Society , 501–522.

Hornbeck, R., Keskin, P., 2014. The historically evolving impact of the ogallala aquifer: Agricultural adaptation to groundwater and drought. American Economic Journal: Applied Economics 6, 190–219.

Houthakker, H.S., 1955. The pareto distribution and the cobb-douglas production function in activity analysis. The Review of Economic Studies 23, 27–31.

Lewis, W.A., 1951. Principles Of Economic Planning. Dennis Dobson Ltd.; London.

Maredia, M.K., Byerlee, D.R., 2000. Efficiency of research investments in the presence of international spillovers: wheat research in developing countries. Agricultural Economics of Agricultural Economists 22.

Matsuyama, K., 1992. Agricultural productivity, comparative advantage, and economic growth. Journal of economic theory 58, 317–334.

Mitchell, B.R., 1982. International historical statistics: Africa and asia .

Moscona, J., 2019. Agricultural development and structural change, within and across countries. Unpublished Manuscript .

Nunn, N., Qian, N., 2011. The potato's contribution to population and urbanization: Evidence from a historical experiment. The Quarterly Journal of Economics 126, 593–650.

Nurkse, R., 1953. Problems of capital formation in underdeveloped countries. Oxford University Press.

Pandey, S., Velasco, M.L., Yamano, T., Walker, T., Alwang, J., et al., 2015. 13 scientific strength in rice improvement programmes, varietal outputs and adoption of improved varieties in south asia. Crop improvement, adoption, and impact of improved varieties in food crops in sub-Saharan Africa , 239–264.

Perez, N., Rosegrant, M.W., 2015. The impact of investment in agricultural research and development and agricultural productivity. IFPRI Discussion Paper 1447 .

Perkins, J.H., 1997. Geopolitics and the green revolution. Oxford University Press.

Pingali, P., Kelley, T., 2007. The Role of International Agricultural Research in Contributing to Global Food Security and Poverty Alleviation: The case of the CGIAR. Elsevier. volume 3 of *Handbook of Agricultural Economics*. chapter 45. pp. 2381–2418.

Pingali, P.L., 2012. Green revolution: Impacts, limits, and the path ahead. Proceedings of the National Academy of Sciences 109, 12302–12308.

Raitzer, D.A., Kelley, T.G., 2008. Benefit-cost meta-analysis of investment in the International Agricultural Research Centers of the CGIAR. Agricultural Systems 96, 108–123.

Renkow, M., Byerlee, D., 2010. The impacts of {CGIAR} research: A review of recent evidence. Food Policy 35, 391 – 402.

Restuccia, D., Yang, D., Zhu, X., 2008. Agriculture and aggregate productivity. Journal of Monetary Economics 55, 234–250.

Roodman, D., Nielsen, M.Ø., MacKinnon, J.G., Webb, M.D., 2019. Fast and wild: Bootstrap inference in stata using boottest. The Stata Journal 19, 4–60.

Rose, B., 1985. Appendix to: The rice economy of Asia, Rice Statistics by Country. The Johns Hopkins University Press.

Rosenstein-Rodan, P.N., 1943. Problems of industrialisation of eastern and south-eastern europe. Economic Journal 53, 202–211.

Rusike, J., Mahungu, N., Jumbo, S., Sandifolo, V., Malindi, G., 2010. Estimating impact of cassava research for development approach on productivity, uptake and food security in Malawi. Food Policy 35, 98–111.

Thirtle, C., Lin, L., Piesse, J., 2003. The Impact of Research-Led Agricultural Productivity Growth on Poverty Reduction in Africa, Asia and Latin America. World Development 31, 1959–1975.

Vollrath, D., 2011. The agricultural basis of comparative development. Journal of Economic Growth 16, 343–370.

von der Goltz, J., Dar, A., Fishman, R., Mueller, N.D., Barnwal, P., McCord, G.C., 2020. Health impacts of the green revolution: Evidence from 600,000 births across the developing world. Journal of Health Economics 74, 102373.

Walker, T.S., Alwang, J., 2015. Crop Improvement, Adoption and Impact of Improved Varieties in Food Crops in Sub-Saharan Africa. CABI.

Figure Captions and Tables

Figure 1

Notes: *This figure shows 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 1965/1966 when wheat HYVs were released in India and a number of other countries in Asia (the dashed vertical line).*

Figure 2

Notes: *This figure shows our baseline crop-level event-study estimates based on Equation 10. The dependent variable is* $\ln yields$*. The sample period is 1945-2010. In Panel A, we show estimates under the assumption that the treatment effect is of the same magnitude across crops. In Panels B-D, we report effects by crop type (cereals, roots and tubers, pulses); all from the same regression. Both regressions (i.e., Panel A and Panels B-D) include controls for:* $\ln harvestedarea$*, country-by-year fixed effects, and crop-by-country fixed effects. The omitted comparison event-year is the year prior to an HYV release (the dashed vertical line). The dashed lines are the 95-percent confidence bands. Standard errors are clustered at the crop-by-country level.*

Figure 3

Notes: *This figure shows the most important food crop for each country, measured at production values in 1966, (Panel A), and predicted GR yields in 2010, 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.*

Figure 4

Notes: *This figure shows our 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. Panels B and D additionally control for pre-Green Revolution income and population growth (1950-1963) interacted with year fixed effects. The sample period is 1950-2010, and the samples are balanced with 85 countries. The dashed vertical line indicate 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

Table 1: Release year of first HYV by crop

Crop	Crop type	IARC research center (host country)	Date of center founding	Date of initial crop mandate	Date of 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

Notes: *This table lists, by crop, the year in which the first HYV was released along with the International Agricultural Research Center (IARC) from which it originated. Our empirical analysis does not include potatoes, soybeans, and plantain for reasons explained in footnote 15. For further details, see Online Appendix Table A1.*

Table 2: Estimated effect on crop yields

	Baseline for all crops	Baseline by crop type	High quality yield data	FAOSTAT yield data	Excluding IARC countries	Excluding pure control crops	Crop-type linear trends
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
HYV x t	0.013***						
	(0.001)						
HYV x t, cereals		0.013***	0.015***	0.012***	0.012***	0.014***	0.017***
		(0.001)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)
HYV x t, roots/tubers		0.009***	0.010***	0.009***	0.009***	0.012***	0.010**
		(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
HYV x t, pulses		0.007***	0.009***	0.007***	0.006***	0.009***	0.007***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
ln harvested area	-0.023**	-0.022**	-0.009	-0.028**	-0.025**	-0.021*	-0.023**
	(0.010)	(0.010)	(0.016)	(0.011)	(0.011)	(0.011)	(0.010)
Country x year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country x 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

Notes: *This table reports estimates based on Equation 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 ($HYV \times t$) in Column 1. This variable is further interacted with crop-type indicator (cereals, roots and tubers, pulses) in Columns 2-7. Columns 1 and 2 constitute the baseline specification. In Column 3, we use only high quality data. In Column 4, only FAOSTAT data are used; so the sample period is 1961-2010. In the remaining columns, the sample period is 1945-2010. In Column 5, we exclude from the sample countries hosting an international agricultural research center. In Column 6, pure control crops are excluded, such that we only include crops for which HYVs became available within the sample period. Standard errors are clustered at the crop-by-country level.* $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

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)
	Panel 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
	Panel 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
	Panel 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
	Panel 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
	Panel 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
	Panel 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

Notes: *This table reports country-level estimates based on Equation 13. Food crops (in Panel B) is defined as 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.* 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, expect for Panel D where the sample ends in 2000. All regressions include country and year fixed effects. Column 1 is the baseline specification. In Column 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 non-parametric controls for temperature and precipitation. Column 4 controls for convergence in the agricultural sector by including* ln *yields in 1961 interacted with year fixed effects. Column 5 excludes countries hosting an International Agricultural Research Center (IARC). Standard errors are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.* $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4: Effects on income and population

	Simple model	Interaction model (baseline)	Controlling for pre-GR income and population growth	Caloric aggregation	Controlling for climate	Excluding IARC host countries
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: *Income*					
$\ln \widehat{GR}$	2.10***					
	(0.56)					
$\ln \widehat{GR} \times initial\ AES$		2.75***	2.69***	2.60***	2.71***	2.78***
		(0.87)	(0.88)	(0.85)	(0.83)	(0.90)
	Panel B: *ln population size*					
$\ln \widehat{GR}$	-0.55***					
	(0.210)					
$\ln \widehat{GR} \times initial\ AES$		-0.600**	-0.84***	-0.53**	-0.60**	-0.63**
		(0.26)	(0.23)	(0.25)	(0.23)	(0.27)
	Panel C: *ln GDP per working age population*					
$\ln \widehat{GR}$	1.61***					
	(0.50)					
$\ln \widehat{GR} \times initial\ AES$		2.25***	2.27***	2.09***	2.21***	2.29***
		(0.79)	(0.81)	(0.78)	(0.76)	(0.83)
Controls (x 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

Notes: *This table reports country-level estimates based on Equations 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. 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: Column 3 controls for pre-Green Revolution growth in GDP per capita and population growth (1950-1963), both interacted with year fixed effects. In Column 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 non-parametric controls for temperature and precipitation. Column 6 excludes countries hosting an International Agricultural Research Center (IARC). Standard errors are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.* $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: Effects on demography and schooling

	Dependent variable:					
	ln adult mortality rate	*ln infant mortality rate*	*ln total fertility rate*	*rate of natural increase*	*net migration rate*	*years of schooling, age 15-20*
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln \widehat{GR} \times initial\ AES$	-1.03*	-2.12***	-2.21***	-0.036***	-0.003	4.07*
	(0.56)	(0.59)	(0.46)	(0.011)	(0.008)	(2.32)
Controls (x 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

Notes: *This table reports country-level estimates based on Equation 14. The rate of natural increase (in Column 4) is calculated as the crude birth rate minus the crude death rate. The net migration rate (in Column 5) is calculated as the population growth rate minus the rate of natural increase. All regressions include country and year fixed effects. 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 are based on a two-step wild-cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.* $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6: Counterfactual scenarios

	Actual	Delayed 10 years	Delayed 25 years	No GR	Delayed 10 years	Delayed 25 years	No GR
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Levels			*Loss/gain compared to actual*		
Yield in 2010 (2010 USD/hectare)	812	687	541	416	-15 %	-33 %	-49 %
Yield growth p.a. 1965-2010 (percent)	2.0	1.6	1.1	0.5	-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 p.a. 1965-2010 (percent)	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 %
Popuation growth p.a. 1965-2010 (percent)	1.9	2.0	2.2	2.3	5 %	11 %	19 %
Popuation 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 %

Notes: *This table shows the effect of the Green Revolution 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: i) a 10-year delay in arrival of the Green Revolution, ii) a 25-year delay in the arrival of the Green Revolution, and iii) 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 tertile is the tertile 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.*

Figure

Figure 1: Wheat yield and date of first HYVs

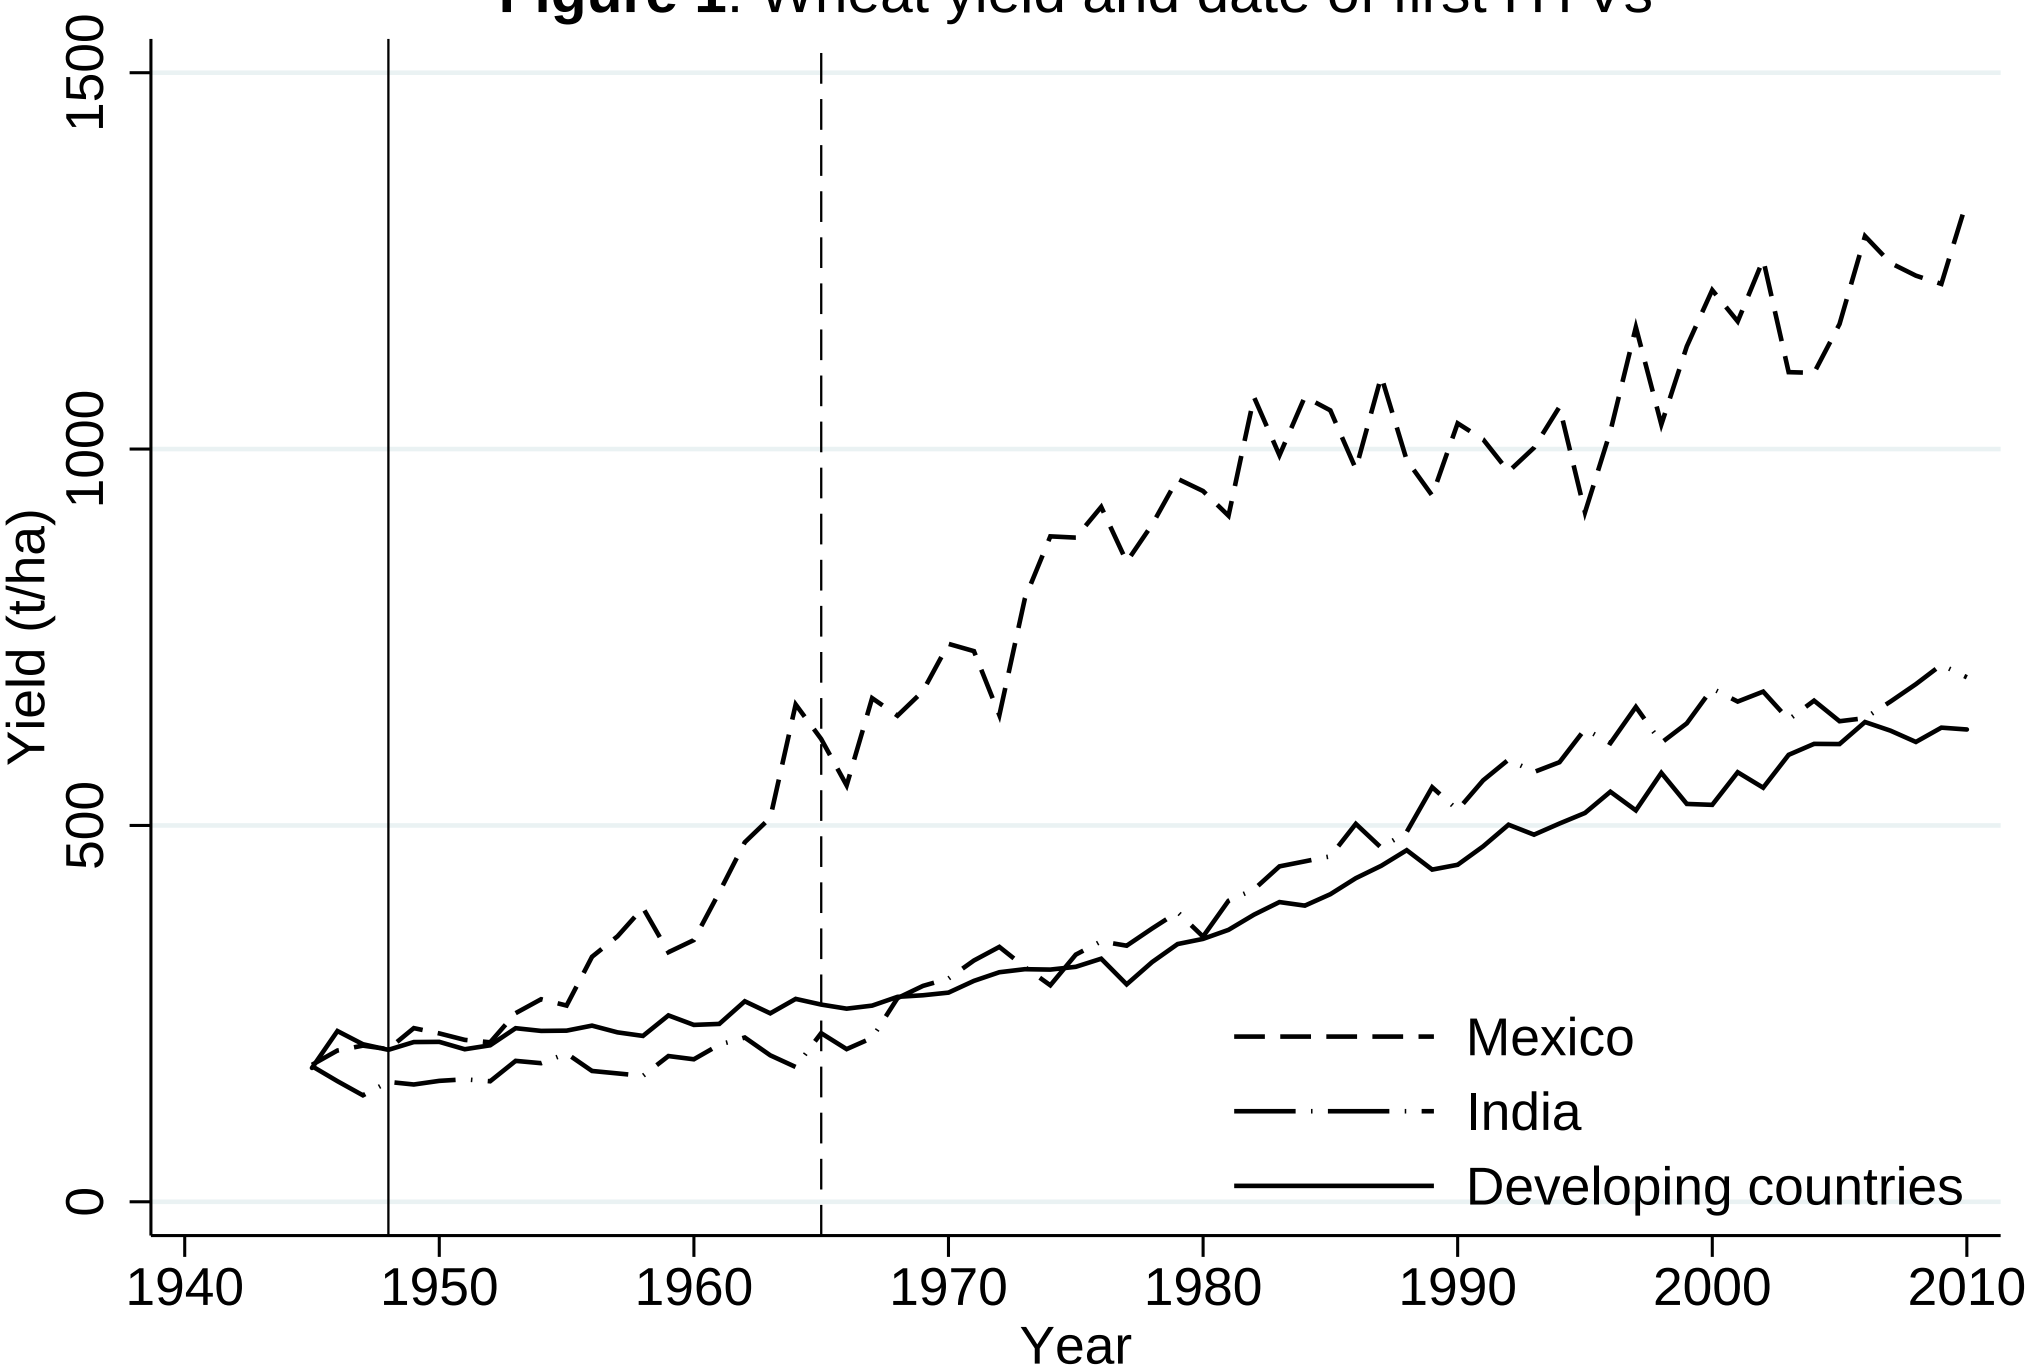

Figure

Figure 2: Event studies for all crops and by crop type

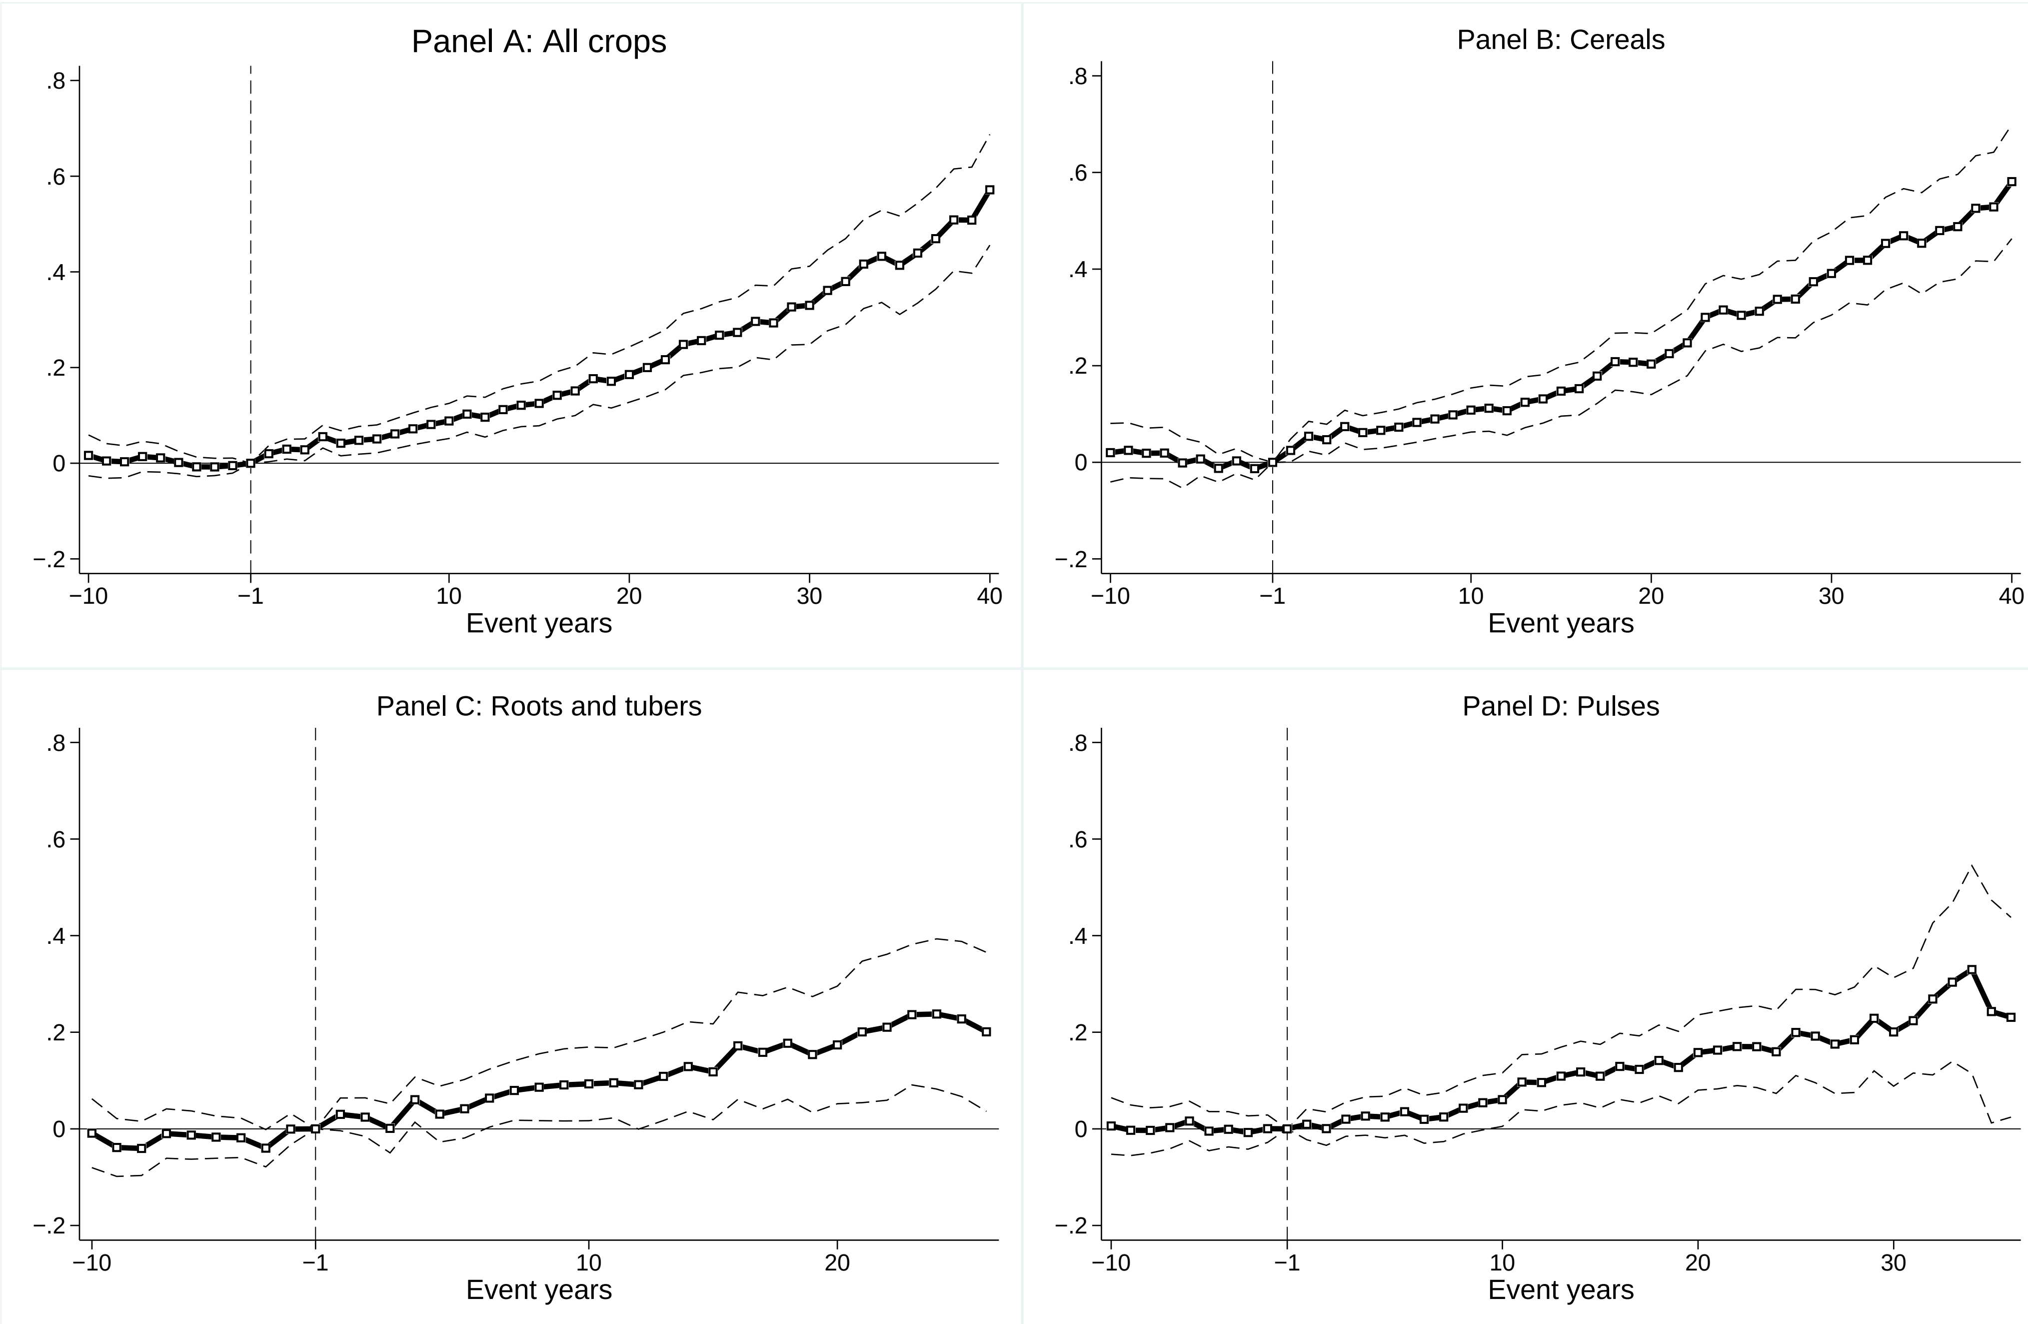

Figure 3: Spatial distribution of the Green Revolution

Panel A: Most important food crops before the Green Revolution

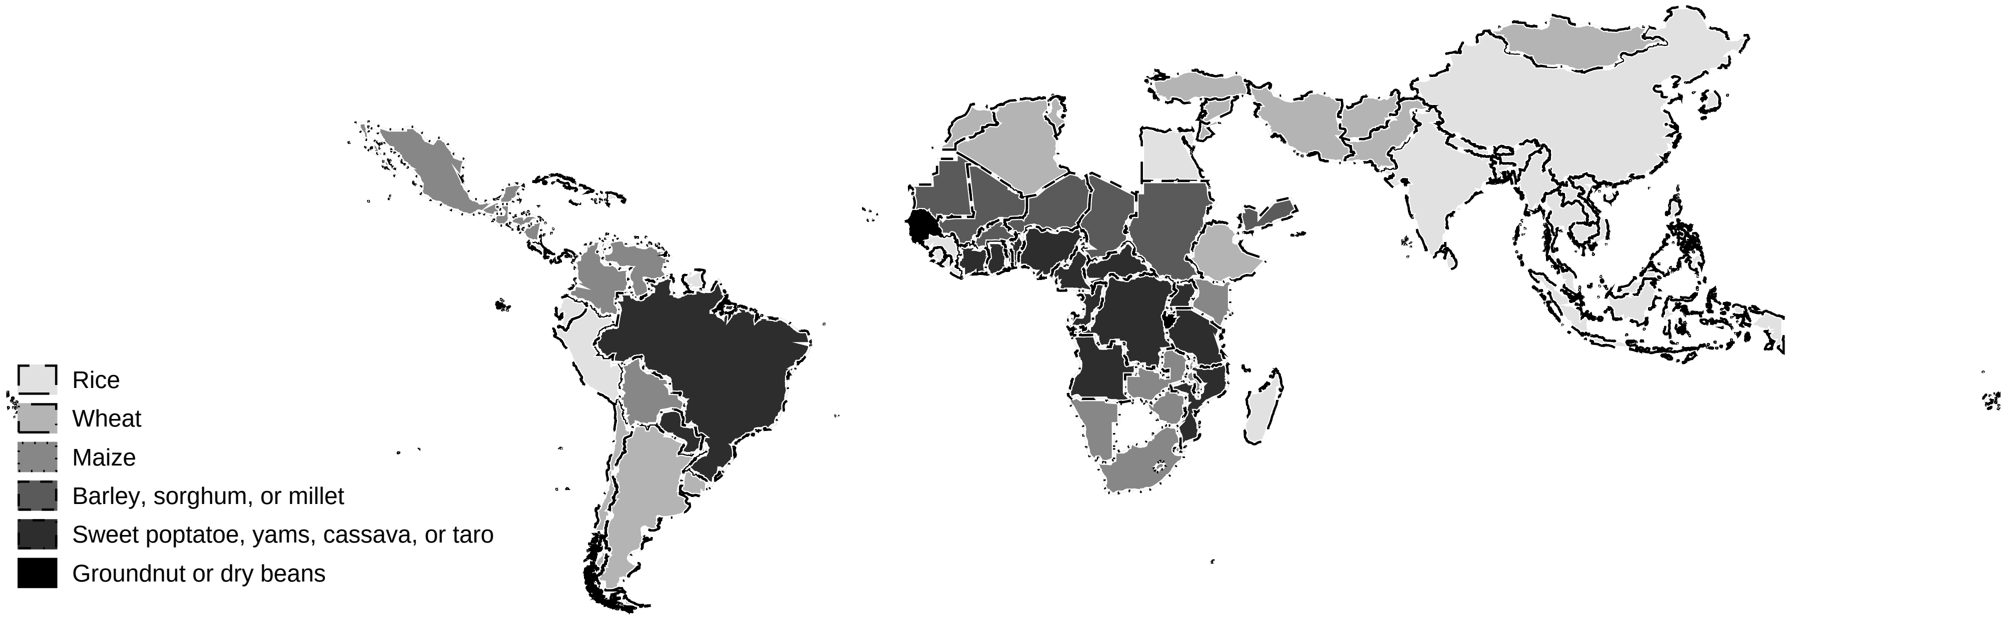

Panel B: Predicted GR yields in 2010

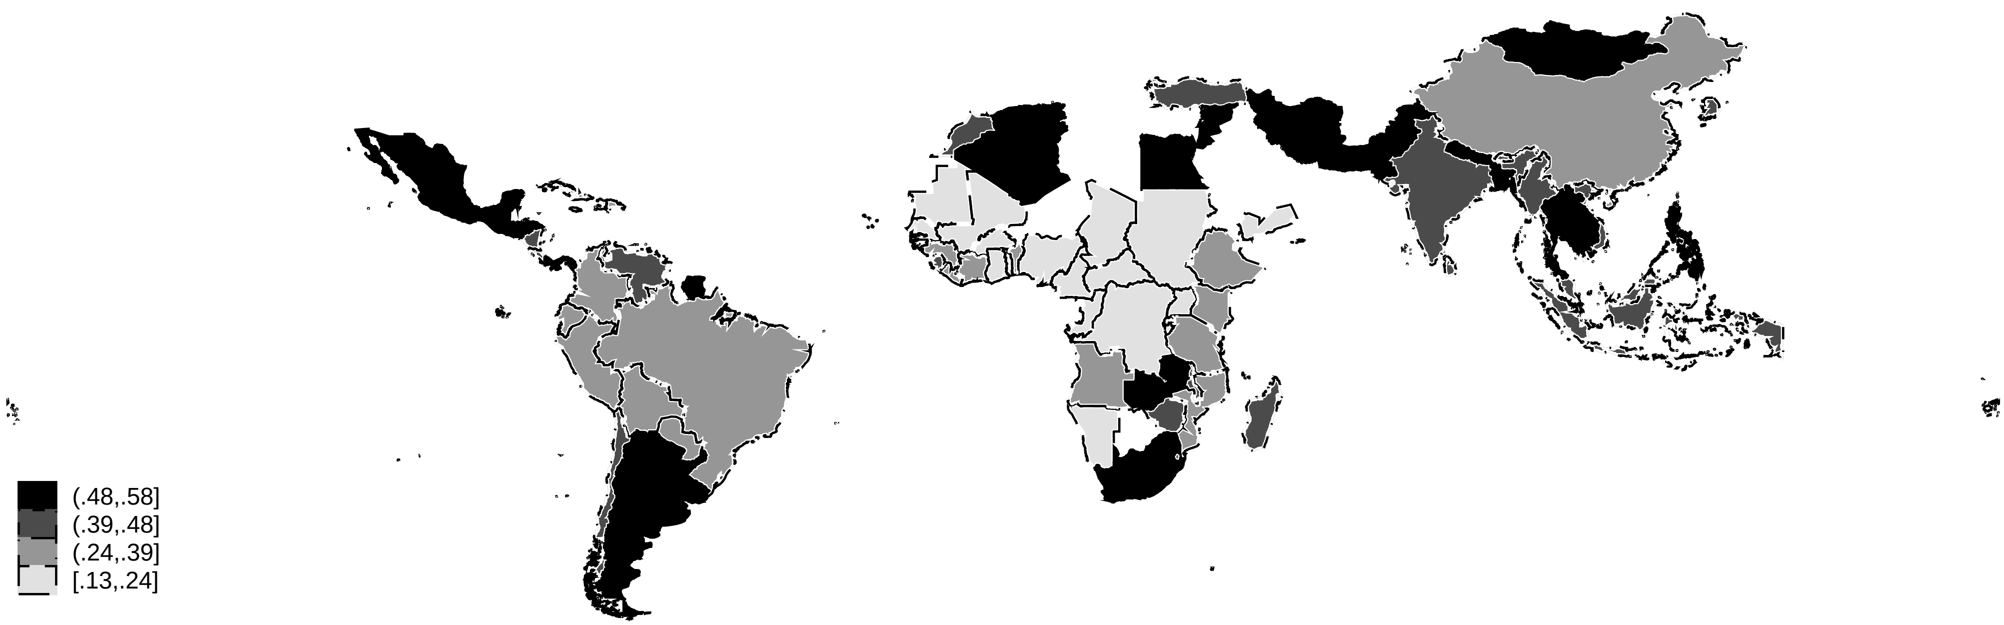

Figure

Figure 4: Event−study estimates for GDP per capita and population

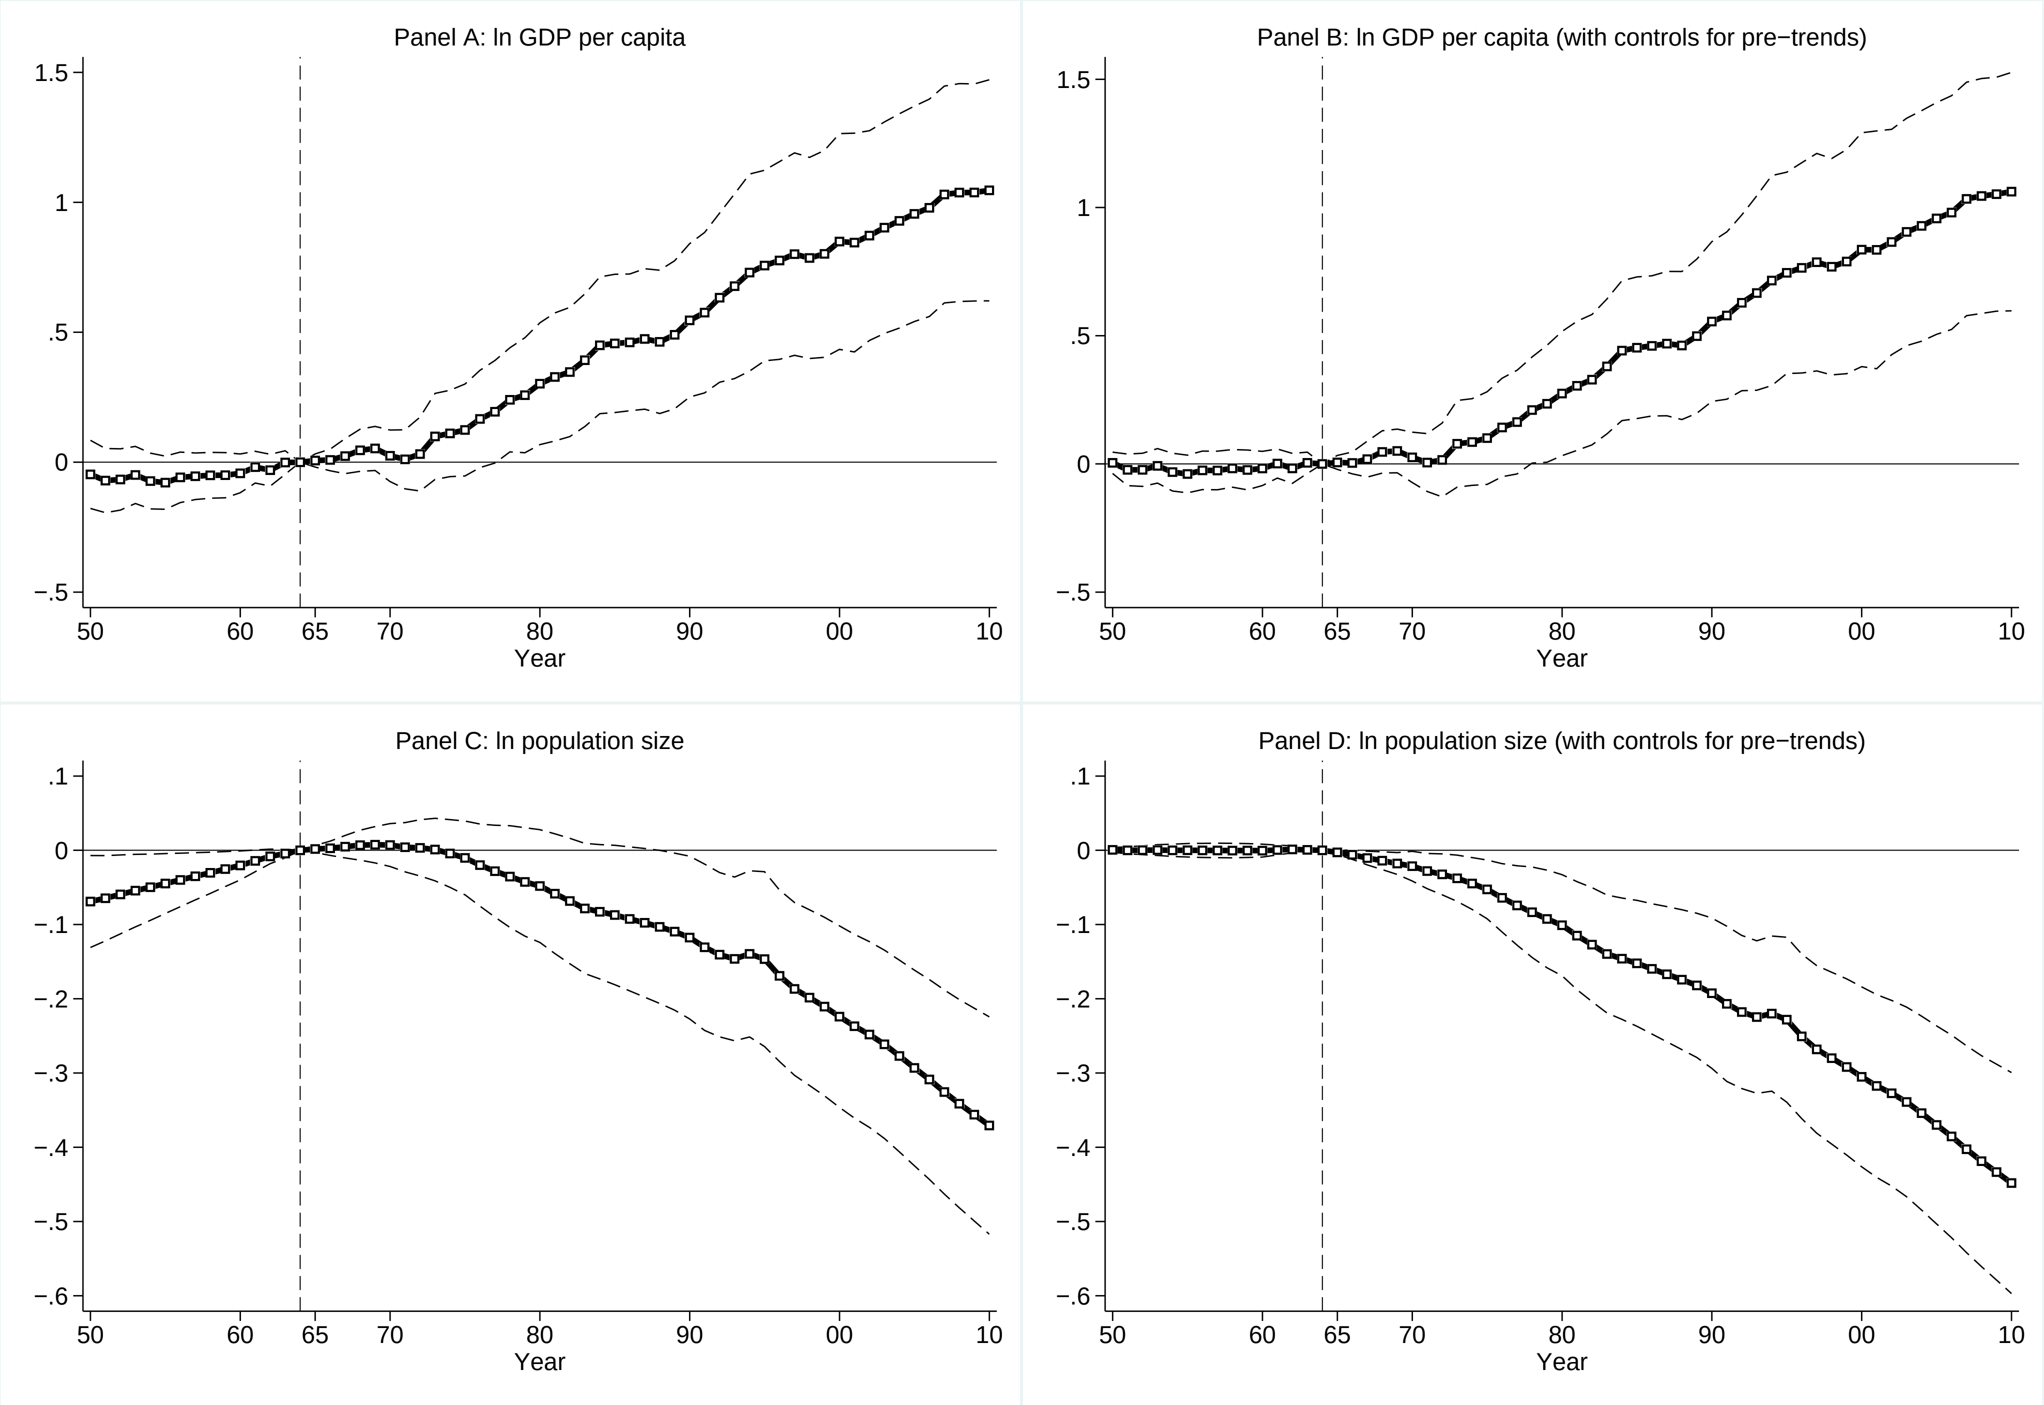

Online Appendix to "Two Blades of Grass: The Impact of the Green Revolution"

Douglas Gollin   Casper Worm Hansen   Asger Mose Wingender*

Abstract

This online appendix contains supporting material to the paper "Two Blades of Grass: The Impact of the Green Revolution". The structure of the appendix follows the structure of the paper. The supporting material is for online publication only.

*Gollin: Oxford University and CEPR, douglas.gollin@qeh.ox.ac.uk; Hansen: University of Copenhagen, casper.worm.hansen@econ.ku.dk; Wingender: University of Copenhagen, amw@econ.ku.dk.

Contents

A Background

This section reports supporting material to Section 2 of the main paper. Appendix Table A1 is an extended version of Table 1 in the paper and provides further details about the development of high-yielding varieties (HYVs) of different crops. Column 1 reports the crops of which HYVs have been released by the international agricultural research centers (IARCs). Specifically, our criterion is the first date at which a variety bred within the institution was made available as an essentially finished product to national programs and partner institutions. This criterion is generally straightforward to interpret, but some cases give rise to slight ambiguity. Our criteria are designed to capture the fact that the underlying institutional innovation was that the new research centers had the ability to conduct cross-fertilization at scale. We thus define the initiation of the Green Revolution in each crop as the date at which novel germplasm was created by scientists and made available for cultivation. Thus, we have sought to exclude material circulated to researchers in early-stage experimental form, as distinct from finished material. We have also sought to restrict our focus to material that required some innovative genetic input, such as cross-fertilization; this means that we omit cases where centers simply took part in the refinement or exchange of existing varieties. Finally, we define ?release? here in a de facto sense rather than a de jure sense; after the 1970s, most of the international centers stopped formally ?releasing? HYVs, since this function seemed to usurp the prerogatives of national programs.

Column 2 lists the names and locations of the IARCs that developed and released the first HYVs of the crops in Column 1. In Columns 3-5, we provide information about the year in which the IARC was founded, the date of the mandate to develop HYVs of the crop, and the year in which the first HYV was released. We use the latter in our empirical analysis. The remaining columns report additional notes and data sources, respectively.

Appendix Table A1: Background on the development of HYVs

Crop	CGIAR research centres	Date of centre founding	Date of initial crop mandate	Date of first released material	Notes on IARC contributions to varietal improvement.	Sources	Year used in baseline analysis	# of HYVs released by IARCs and NARS as of today
Barley	ICARDA (Syria; currently displaced to Morocco and Jordan)	1977	1977	1979	ICARDA benefited from ongoing research at the time of its founding in 1977; some research was done by precursor breeding programs, but the first two varieties released after the founding of the center were in 1979, with Holkr and Karoon, released in Ethiopia and Iran respectively.	ICARDA varietal release data base: http://oldindms.codeobia.com/variety-release	1979	246 (as of 2017)
Dry beans	CIAT (Colombia)	1970	1973	1979	Some breeding happened already in the 1930s and 1940s. The first CIAT-developed releases came in the late 1970s. First releases in Africa were in 1996 (according to CIAT website), although the list of released varieties suggests that there were already some selections of traditional varieties that were being released in Africa in the 1980s.	From CIAT list of released varieties	1979	> 300 (as of 2000)
Cassava	CIAT, IITA (Colombia, Nigeria respectively)	1970	1973	1981 for LAC, 1984 for Asia, 1984 for Africa	IITA pursued an ambitious agenda focused initially on selecting existing varieties with resistance to cassava mosaic virus (CMV). Some varieties were tested as early as 1976, but TMS 30572 was the first variety released by IITA -- known as Nicass 1 in Nigeria, the largest global producer of cassava. Release was in 1984. The first CIAT releases in Latin America are dated to 1981 and 1984 for Asia.	From CIAT list of released varieties; IITA data produced for DIIVA.	1984	>400 by IITA + 22 CIAT-developed varieties through 1997
Chickpea	ICARDA (Syria; currently displaced to Morocco and Jordan)	1977	1977	1984	Breeding programme initiated 1977/78. The variety Yialousa (ILC 3279) was released in 1984 in Cyprus and is the first release claimed by ICARDA.	ICARDA varietal release data base: http://oldindms.codeobia.com/variety-release	1984	169 (as of 2017)
Cowpea	IITA (Nigeria)	1967	Approximately 1970. (Early work at IITA was focused on eco-regional mandates, rather than crop mandates.)	1974	IITA began work by screening and testing traditional varieties, with hybridization work taking place in the mid 1970s. The first released variety based on IITA breeding appears to be VITA-5, released in Togo in 1974.	IITA (2017), *IITA at 50*; DIIVA data; Ortiz (1998); Redden et al. (1984).	1974	200 (as of 2011)
Faba bean (broad bean)	ICARDA (Syria; currently displaced to Morocco and Jordan)	1977	1977	1986	The variety Barkat (ILB 1269) was released in 1986 in Iran, and ICARDA identifies this as the first released faba variety from their research.	ICARDA varietal release data base: http://oldindms.codeobia.com/variety-release	1986	73 (as of 2017)
Groundnut	ICRISAT (India)	1972	1974?	1985	ICGS 11 was recommended for release in 1985. It was the first ICRISAT-bred groundnut variety released at the All-India level.	T.S. Walker, personal communication (2019)	1985	NA
Lentils	ICARDA (Syria; currently displaced to Morocco and Jordan)	1977	1977	1980	The first lentil varieties released were in Ethiopia in 1980, R186 and EL142 (ILL5071).	ICARDA varietal release data base: http://oldindms.codeobia.com/variety-release	1980	154 (as of 2017)
Maize	CIMMYT, IITA (Mexico, Nigeria respectively)	1963	1963	1966 or so.	Hybrid rice technology was developed in the United States and diffused across the U.S. during the middle third of the 20th century. Some early hybrids and open-pollinated varieties were released in developing countries beginning in the 1950s. Early varieties targeted temeprate zones (e.g., in Kenya). Beginning in the 1960s, CIMMYT developed a series of improved populations of maize, based heavily on the so-called Tuxpeno germplasm. These were used to develop both open-pollinated and hybrid varieties, with releases beginning in 1966. Maize varieties developed by CIMMYT and IITA generally targeted tropical zones. A hybrid maize programme at IITA was added in 1979, with the first hybrid maize varieties released in West and Central Africa by 1983.	Byerlee (2016); Byerlee (1994); CIMMYT (1992).	1966	603 (as of 2018)
Millet	ICRISAT (India)	1972	1972	1982	ICRISAT released its first millet variety, WC-C75, in 1982 (1983 in Africa).	T.S. Walker, personal communication (2019)	1982	NA
Pigeon Pea	ICRISAT (India)	1972	1972	1991	First HYV variety released by ICRISAT in 1991	T.S. Walker, personal communication (2019)	1991	NA
Plantain	IITA (Nigeria)	1967	1972 (no research mandate before then)	2002	IITA initiated a research program. Much of the early work consisted of screening germplasm and trying to identify sources of disease resistance. Because banana and plantain are difficult to hybridize, early efforts also focused on making disease-free planting materials available to farmers. There is clear evidence of hybrid varieties released in 2002, but there is no clear evidence of previous releases based on IITA or INIBAP breeding.	IITA (2017), *IITA at 50*.	2002	NA
Potato	CIP (Peru)	1971	1971	1990	Modern varieties of potatoes have been released at least since 1927 (Kenya), but at a slow pace, based on varieties selected for Europe and North America. The pace of breeding accelerated in the 1970s, based on material from a predecessor Rockefeller Foundation program. The variety Canchan-INIA was released in Peru ini 1990 and was the first released variety actually developed at CIP.	T.S. Walker, personal communication (2019)	1990 if we go by a strict definition	NA
Rice	IRRI (Philippines). Later also WARDA (Liberia) for Africa and CIAT (Colombia) for Latin America	1960	1960	1966	Some varietal improvement efforts took place under colonial governments in India and some other Asian countries. An International Rice Commission led by the UN Food and Agriculture Organisation had some success in developing high-yielding tall varieties (Mahsuri and Biplab) in the early 1950s. The first true HYV in rice was the semi-dwarf variety IR8, released in 1966 by IRRI. IRRI subsequently produced a stream of varieties; after 1980 IRRI no longer released varieties directly but only through NARS partners. CIAT released their first variety in 1971. WARDA in West Africa initially did little breeding work, but in recent years, in a new incarnation as AfricaRice, it has developed the Nerica rice lines and carried out other breeding efforts.	Dalrymple (1974, 1985, 1986); Herdt and Capule (1983), Barker et al. (1985)	1966	>1500 in Asia alone
Sorghum	ICRISAT (India)	1972	1972	1983	HYV sorghum breeding started in India in the early 1960s, and some material may have filtered into All-India breeding programs in the 1970s. However, data show that the first variety released that was based on ICRISAT breeding appears to be hybrid Hageen Dura, released in the Sudan in 1983.	T.S. Walker, personal communication (2019)	1970 (or 1983?)	NA
Soybean	IITA (Nigeria)	1967	Unclear	1979	The first released soybean variety that appears to have been developed at IITA is a cross of Jupiter (an old variety originating in the US) with a variety called Anidazo (or possibly Anidaso). Soybean was a mandate crop of IITA since its early years, but there were major technical problems in adapating germplasm from temperate zones so that they would grow effectively in the tropics. Nevertheless, by the mid 1980s, there were considerable nubers of varieties being produced.	IITA (2017); Pachico (2014)	1979	201 (as of 2012)
Sweet Potato	CIP (Peru)	1971	1988	mid-1980s	The mandate for sweet potato research was originally given to IITA, but it passed to CIP in 1988. Up to that point, no useful material had been developed for the relevant areas in Central and East Africa, because of susceptibility to two important viruses. Beginning in 1999, the first NASPOT varieties were released in Uganda by CIP and selected by the National Program in Uganda.	T.S. Walker, personal communication (2019)	1999	NA
Wheat	CIMMYT (Mexici). Later also ICARDA (Syria) for MENA countries	1963	1963	1965-66 (first releases beyond Mexico)	CIMMYT's work had antecedents in a program funded by the Rockefeller Foundation that began work in Mexico in 1943, focused primarily on rust. Work on semi-dwarf varieties began in mid 1950s (Byerlee, *Birth of CIMMYT*). The first HYVs were released in Mexico in 1948, with yields increasing sharply in the 1950s. Beginning in 1963, the CIMMYT high-yielding wheats were taken to India and Pakistan, where they were first released in 1965/66.	Byerlee (2016); Dalrymple (1978); Heisey et al. (2002); Pingali (1999), CIMMYT (1992).	1965	>2,200 by 1997
Yam	IITA (Nigeria)	1967	Approximately 1970. (Early work at IITA was focused on eco-regional mandates, rather than crop mandates.)	1990	IITA has had a research program since soon after its founding, but the earliest variety attributable to IITA breeding is TDr 608, which was released for the first time in 1990.	IITA at 50; DIIVA data and DIIVA narrative.	1990	79 (as of 2010)

Notes: DIIVA data are available at https://www.asti.cgiar.org/diiva. The date of initial crop mandate is the date at which the CGIAR center was first assigned a formal mandate for breeding in that crop. Additional sources: Personal communication, Thomas S. Walker (2019), Derek Byerlee (2019), Rodomiro Ortiz (2019), Abdoul Aziz Niane (2019).

B Theoretical framework

This section contains the model outlined in Section 3 of the paper, including all the intermediate steps needed to derive the equations reported in the paper.

Consider a country with N regions of fixed size $\bar{X}_k$, $k = 1, 2, .., N$. A fraction of land, s_{kt}, within each region is used for crop production, the rest is left fallow. Let $X_{kt} = s_{kt}\bar{X}_k$ denote crop land in region k at time t. To simplify, assume that regions have distinct agro-ecologies such that region k can only grow 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 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 all plots $i > X_{kt}$ are left fallow. Moreover, let a_{ik} distributed smoothly and differentiably across plots in such a way that the average productivity of crop land within a region is given by:

$$\frac{1}{X_{kt}} \int_0^{X_{kt}} a_{ikt} di = a_{kt} s_{kt}^{-\delta},\ 0 < \delta < 1\,, \tag{A1}$$

where a_{kt} is a region-specific (and by implication crop-specific) productivity parameter to be specified below. 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. Land and labor are, in other words, perfect complements at the farm level, such that farm-level production function is Leontief. The number of farms in region k is consequently equal to both aggregate crop land in region k, denoted X_k , and to total farm labor in region k, which we denote L_k. These assumptions, along with Equation A1, imply that aggregate agricultural output in region k is a Cobb-Douglas function of total land and labor in the region:

$$Y_{Akt} = a_{kt} s_{kt}^{-\delta} X_{kt} = a_{kt} s_{kt}^{-\delta} L_{kt} = a_{kt} \bar{X}_k^{\delta} L_{kt}^{1-\delta}. \tag{A2}$$

Total land in region k, $\bar{X}_k$, is fixed, and there are diminishing returns to labor in the aggregate production function because additional farmers pull less fertile land into agricultural use. We abstract from capital, fertilizer, and other inputs. This simplification does not change the qualitative predictions of the model, but including such adjustable inputs would tend to amplify the positive yield shock from HYVs.

The average yield in region k, which is equal to the average yield of crop k, is:

$$yield_{kt} = \frac{Y_{Akt}}{X_{kt}} = a_{kt} s_{kt}^{-\delta}. \tag{A3}$$

To highlight the effect of HYVs, assume that the productivity parameter a_{kt} can be decomposed as:

$$a_{kt} = a_{kt}^{HYV} \bar{a}_k z_t u_{kt} \,, \tag{A4}$$

where $a_{kt}^{HYV} = 1$ if no HYVs are available, and $a_{kt}^{HYV} > 1$ if HYVs are grown. a_{kt}^{HYV} varies from crop to crop to allow for different release dates of the first HYV, and for differential effects on yields. a_{kt}^{HYV} for the individual crops also vary over time to allow for newer, higher-yielding HYVs being released following the first one. The parameter $\bar{a}_k > 0$ is a crop-specific time-invariant productivity level, $z_t > 0$ is a country-wide productivity trend, and $u_{kt} > 0$ is an idiosyncratic productivity shock to crop k with mean 1. We can consequently write:

$$\ln yield_{kt} = \ln a_{kt}^{HYV} - \delta \ln s_{kt} + \ln \bar{a}_k + \ln z_t + \ln u_{kt}. \tag{A5}$$

This equation is the theoretical counterpart of our main crop-level estimating equation.

We now turn to the aggregate implications of introducing HYVs. To simplify notation, we set $\ln z_t = \ln u_{kt} = 0$, thereby assuming that the releases of HYVs are the only thing affecting yield growth. Let $t = 0$ be the period immediately 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:

$$GR_t = \sum_{k=1}^{N} a_{kt}^{HYV} \frac{Y_{Ak0}}{Y_{A0}} \,, \tag{A6}$$

where $\frac{Y_{Ak0}}{Y_{A0}}$ is the share of crop k in aggregate crop production prior to the Green Revolution.

Let Y_{At} denote aggregate agricultural output, and let X_t denote aggregate crop land. Actual aggregate yields (in quantities) are:

$$\begin{aligned} yield_t = \frac{Y_{At}}{X_t} &= \frac{\sum_{k=1}^{k=N} Y_{Akt}}{\sum_{k=1}^{k=N} X_{kt}} \\ &= \frac{\sum_{k=1}^{N} a_{kt} \bar{X}_k^{\delta} L_{kt}^{1-\delta}}{X_t} \\ &= \frac{GR_t}{GR_t} \frac{\sum_{k=1}^{N} a_{kt} \bar{X}_k^{\delta} L_{kt}^{1-\delta}}{X_t} \\ &= \frac{Y_{A0}}{X_0} GR_t \frac{\sum_{k=1}^{N} a_{kt} \bar{X}_k^{\delta} L_{kt}^{1-\delta}}{\sum_{k=1}^{N} a_{kt} \bar{X}_k^{\delta} L_{k0}^{1-\delta}} \frac{X_0}{X_t}, \end{aligned} \tag{A7}$$

where we in the last step use that $a_{kt}^{HYV} \bar{a}_k = a_{kt}$ in the absence of other shocks to productivity. To turn this expression into something that lend itself to estimation, we assume free movement of labor, that unused land is costless, and that farmers maximize their incomes. These assumptions imply that marginal farmers in each region will have the same income. Because land is free, farmers earn the entire output of their farms (this is a common assumption in models of agriculture in developing countries. Free land can also be seen as an analogy of non-market based allocation of land in traditional societies). The incomes of the marginal farmers are consequently equal to the marginal productivity of land in each region. By implication, we have that for two regions j and k:

$$\forall j, k \leq N : a_{jt} \bar{X}_j^{\delta} L_{jt}^{-\delta} = a_{kt} \bar{X}_k^{\delta} L_{kt}^{-\delta}\,, \tag{A8}$$

which can be rearranged to an expression for labor in region j as a function of labor in region k, relative region sizes, and relative productivity levels:

$$L_{jt} = \left(\frac{a_{jt}}{a_{kt}} \right)^{\frac{1}{\delta}} \frac{\bar{X}_j}{\bar{X}_k} L_{kt}. \tag{A9}$$

Sum across regions to obtain:

$$\begin{aligned} L_{At} = \sum_{j=1}^{N} L_{jt} &= \sum_{j=1}^{N} \left(\frac{a_{jt}}{a_{kt}} \right)^{\frac{1}{\delta}} \frac{\bar{X}_j}{\bar{X}_k} L_{kt} \\ \Leftrightarrow L_{kt} &= \frac{\bar{X}_k a_{kt}^{\frac{1}{\delta}}}{\sum_{j=1}^{N} a_{jt}^{\frac{1}{\delta}} \bar{X}_j} L_{At}\,, \end{aligned} \tag{A10}$$

where L_{At} is aggregate agricultural labor. Plug this expression into Equation A7 to get:

$$
\begin{aligned}
yield_t &= \frac{Y_{A0}}{X_0} GR_t \frac{\sum_{k=1}^N a_{kt} \bar{X}_k^\delta \left(\frac{\bar{X}_k a_{kt}^{\frac{1}{\delta}}}{\sum_{j=1}^N a_{jt}^{\frac{1}{\delta}} \bar{X}_j} L_{At} \right)^{1-\delta}}{\sum_{k=1}^N a_{kt} \bar{X}_k^\delta \left(\frac{\bar{X}_k a_{k0}^{\frac{1}{\delta}}}{\sum_{j=1}^N a_{j0}^{\frac{1}{\delta}} \bar{X}_j} L_{A0} \right)^{1-\delta}} \frac{X_0}{X_t} \\
&= \frac{Y_{A0}}{X_0} GR_t \frac{\sum_{k=1}^N a_{kt}^{\frac{1}{\delta}} \bar{X}_k}{\sum_{k=1}^N a_{kt} a_{k0}^{\frac{1-\delta}{\delta}} \bar{X}_k} \left(\frac{\sum_{j=1}^N a_{j0}^{\frac{1}{\delta}} \bar{X}_j}{\sum_{j=1}^N a_{jt}^{\frac{1}{\delta}} \bar{X}_j} \right)^{1-\delta} \frac{L_{At}^{1-\delta}}{L_{A0}^{1-\delta}} \frac{X_0}{X_t}. \qquad \text{(A11)}
\end{aligned}
$$

Using that $\sum_{j=1}^N a_{jt}^{\frac{1}{\delta}} \bar{X}_j = \sum_{k=1}^N a_{kt}^{\frac{1}{\delta}} \bar{X}_k$ and $L_{At} = X_t$, we get:

$$
\begin{aligned}
yield_t &= \frac{Y_{A0}}{X_0} GR_t \frac{\sum_{k=1}^N a_{kt}^{\frac{1}{\delta}} \bar{X}_k}{\sum_{k=1}^N a_{kt} a_{k0}^{\frac{1-\delta}{\delta}} \bar{X}_k} \left(\frac{\sum_{k=1}^N a_{k0}^{\frac{1}{\delta}} \bar{X}_k}{\sum_{k=1}^N a_{kt}^{\frac{1}{\delta}} \bar{X}_k} \right)^{1-\delta} \left(\frac{X_t}{X_0} \right)^{-\delta} \\
&= \frac{Y_{A0}}{X_0} GR_t \frac{\left(\sum_{k=1}^N a_{kt}^{\frac{1}{\delta}} \bar{X}_k \right)^\delta}{\sum_{k=1}^N a_{kt} a_{k0}^{\frac{1-\delta}{\delta}} \bar{X}_k} \left(\sum_{k=1}^N a_{k0}^{\frac{1}{\delta}} \bar{X}_k \right)^{1-\delta} \left(\frac{X_t}{X_0} \right)^{-\delta} \\
&= \frac{Y_{A0}}{X_0} GR_t \frac{\left(\sum_{k=1}^N \left(a_{kt}^{HYV} \bar{a}_k \right)^{\frac{1}{\delta}} \bar{X}_k \right)^\delta}{\sum_{k=1}^N a_{kt}^{HYV} \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k} \left(\sum_{k=1}^N \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k \right)^{1-\delta} \left(\frac{X_t}{X_0} \right)^{-\delta}. \qquad \text{(A12)}
\end{aligned}
$$

In the last step, we use that $a_{k0}^{HYV} \equiv 1$. To simplify notation in the main paper, we write this expression for yields as:

$$
yield_t = \frac{Y_{A0}}{X_0} GR_t \Phi_t \left(\sum_{k=1}^N \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k \right)^{1-\delta} \left(\frac{X_t}{X_0} \right)^{-\delta}, \qquad \text{(A13)}
$$

where:

$$
\Phi_t \equiv \frac{\left(\sum_{k=1}^N \left(a_{kt}^{HYV} \bar{a}_k \right)^{\frac{1}{\delta}} \bar{X}_k \right)^\delta}{\sum_{k=1}^N a_{kt}^{HYV} \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k}. \qquad \text{(A14)}
$$

Because $GR_0 = \Phi_0 = 1$, we can now write yield growth between period 0, i.e., before the Green Revolution, and a period t after the Green Revolution as:

$$
\Delta \ln yield_t = \ln GR_t + \ln \Phi_t - \delta \Delta \ln X_t. \qquad \text{(A15)}
$$

To pin down the level of agricultural production, assume that households require a subsistence level of food, denoted $\bar{c}_a$, to survive. When that level is achieved, no more food is demanded. This approach to modeling subsistence food requirements has been used by Gollin, Parente, and Rogerson 2007, among others. 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 $yield_t \cdot X_t$. Combined with Equation A13, this equilibrium condition implies that:

$$\begin{aligned} \bar{c}_a L_t &= yield_t \cdot X_t \\ \Leftrightarrow \bar{c}_a L_t &= \frac{Y_{A0}}{X_0} GR_t \Phi_t \left(\sum_{k=1}^{N} \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k \right)^{1-\delta} \left(\frac{X_t}{X_0} \right)^{-\delta} X_t \\ \Leftrightarrow X_t &= \left(\frac{\bar{c}_a X_0^{1-\delta}}{Y_{A0}} \frac{L_t}{GR_t \Phi_t \left(\sum_{k=1}^{N} \bar{a}_k^{\frac{1}{\delta}} \bar{X}_k \right)^{1-\delta}} \right)^{\frac{1}{1-\delta}} . \end{aligned} \tag{A16}$$

Inserting this expression into Equation A15 and taking logs gives us:

$$\Delta \ln yield_t = \frac{1}{1-\delta} \ln GR_t + \frac{1}{1-\delta} \ln \Phi_t - \frac{\delta}{1-\delta} \Delta \ln L_t . \tag{A17}$$

Additionally, from Equation A16 and the fact that $X_t = L_{At}$, we obtain:

$$\Delta \ln L_{At} = \Delta \ln X_t = -\frac{1}{1-\delta} \ln GR_t - \frac{1}{1-\delta} \ln \Phi_t + \frac{1}{1-\delta} \Delta \ln L_t . \tag{A18}$$

By subtracting $\Delta \ln L_t$ on both sides, we get:

$$\Delta \ln \frac{L_{At}}{L_t} = -\frac{1}{1-\delta} \ln GR_t - \frac{1}{1-\delta} \ln \Phi_t + \frac{\delta}{1-\delta} \Delta \ln L_t = -\Delta \ln yield_t . \tag{A19}$$

Suppose now that there in addition to agriculture is a manufacturing sector producing output according to $Y_{Mt} = mL_{Mt}$, where m is a constant productivity term and L_{Mt} denotes the labor force employed in manufacturing. 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:

$$y_t = \frac{Y_t}{L_t} = \frac{Y_{At}}{L_t} + \frac{Y_{Mt}}{L_t} = \bar{c}_a + p_t m \frac{L_{Mt}}{L_t} , \tag{A20}$$

where p_t is the relative price of manufactured goods. For the purpose of our empirical application, it is convenient to focus on growth in GDP per capita (in period 0 prices), which under the model can be approximated as:

$$\begin{aligned}\Delta \ln y_t &\approx \frac{y_t - y_0}{y_0} = \frac{\bar{c}_a + p_0 m \frac{L_{Mt}}{L_t} - \left(\bar{c}_a + p_0 m \frac{L_{M0}}{L_0}\right)}{y_0} \\ &= \frac{p_0 m \Delta \frac{L_{Mt}}{L_t}}{y_0} = -\frac{p_0 m \Delta \frac{L_{At}}{L_t}}{y_0} \\ &\approx -\frac{p_0 m \Delta \ln \frac{L_{At}}{L_t}}{y_0} \frac{L_{A0}}{L_0} \\ &= \frac{p_0 m}{y_0} \frac{L_{A0}}{L_0} \Delta \ln yield\,. \end{aligned} \tag{A21}$$

A shock to yields is not only moderated by the size of the labor force initially employed in agriculture, but also by the term $\frac{p_0 m}{y_0}$, which measures how much more productive workers are in manufacturing than in the economy as such. If the manufacturing sector is more productive than agriculture, $\frac{p_0 m}{y_0}$ would be larger than one, and reallocation of labor to manufacturing would increase GDP. In our theoretical setting, a productivity advantage of manufacturing can be rationalized if manufacturing is located in cities with higher cost if living, if individuals get utility from living in rural areas, or if there are barriers to labor mobility.

To get a sense of how important yields are for growth in developing countries, it is useful to put some numbers on the two fractions on the right-hand side of Equation A21 . We do not have pre-Green Revolution data on $\frac{p_0 m}{y_0}$ for most of the countries we study, but contemporary estimates show that labor productivity in developing countries is higher in nonagriculture than in agriculture, with estimates ranging from three times larger to ten times larger (see Gollin, Lagakos, and Waugh (2014)). We recognize that recent work by Hicks et al. (2017) and **?**, among others, suggest that sectoral productivity differences are smaller when controlling for individual worker characteristics; however, for the purpose of our back-of-the-envelope calculation, we assume that average sectoral productivity is close to marginal productivity. Assuming that labor productivity is three times higher in nonagriculture, and that $\frac{L_{A0}}{L_0} = 0.7$, corresponding to the agricultural employment share in the average developing country before the Green Revolution, implies that $\frac{p_0 m}{y^0} = 1.9$. Under these assumptions, the elasticity of GDP per capita with respect to yields is $1.9 \cdot 0.7 = 1.3$, meaning that a yield shock is amplified by the resulting structural transformation to such an extent that the effect on GDP per capita is larger than the isolated effect in agriculture.

C Crop-level DD estimation

This section explains why our differences-in-differences (DD) design allows us to estimate HYVs' contribution to crop yields. Our theoretical framework leads to the following crop-level estimation equation (Equation 4 in the main paper, and Equation A5 above):

$$\ln yield_{kt} = \ln a_{kt}^{HYV} - \delta \ln s_{kt} + \ln \bar{a}_k + \ln z_t + \ln u_{kt}, \tag{A22}$$

We do not observe $\ln a_{kt}^{HYV}$ in our data, but DD estimation allows us to estimate it. For simplicity, assume that there is only one pre-Green Revolution year ($t = 0$) and only one post-Green Revolution year ($t = 1$). In addition, we have one treated crop ($k = 1$), in which a HYV becomes available in period 1, and one untreated control crop ($k = 2$). We additionally assume that $\delta \ln s_{kt}$ is constant, or equivalently, that we control for it in our regressions. By comparing the relative yields of these crops before and after the Green Revolution, we obtain the following simple DD estimate:

$$DD = (\ln yield_{11} - \ln yield_{10}) - (\ln yield_{21} - \ln yield_{20}) = \ln a_{k1}^{HYV} - \ln a_{k0}^{HYV}, \tag{A23}$$

Because HYVs did not exist before the Green Revolution, we have that $\ln a_{k0}^{HYV} = 0$, and the DD estimate simplifies to:

$$DD = \ln a_{k1}^{HYV}, \tag{A24}$$

which (approximately) is the percentage increase in yields of crop 1 due to HYVs. Intuitively, comparing the crop to itself before and after (i.e., the first "diff") eliminates crop-specific effects, while comparing the treated crop to the untreated crop (i.e., the second "diff") eliminates common time trends. The only thing that remains is the effect of HYVs.

D Data sources and definitions

The empirical analysis of the main paper has two different levels of aggregation: the crop-level analysis and the country-level. This section provides data sources and variables definitions according to these two levels of aggregation.

D.1 Crop level data

D.1.1 Agricultural production and yields 1961-2010

We obtain annual data on total production, harvested area, and yields for individual crops from FAOSTAT ("*Food and Agriculture Organization of the United Nations; FAO*"). Total production for a crop is of course the product of yield and area harvested; in practice, countries report all three, but they may only measure two directly. Some countries measure all three directly and triangulate to produce internally consistent data. Practices vary widely, as does the reliability of the data. The FAO data cover 158 different crops, including all crops for which HYV became available (see the list in Appendix Table A1), the pure control crops in our analysis (bambara beans, buckwheat, canary seed, fonio, lupins, quinoa, taro, vetches, and yautia), and a wide range of different cash crops. The classification of crops into types (e.g., cereals and pulses) we use in the paper is also taken from FAO.

D.1.2 Agricultural production and yields 1920-1960

We have collected historical data on production, harvested area, and yields in order to test for pre-Green Revolutions trends at the crop level over a longer period than the FAO data allow us. The resulting data set has annual observations from 1920 to 1961 for barley, cassava, dry beans, lentils, maize, millet, pigeonpeas, rice sorghum, sweet potatoes, and wheat. The coverage is more complete for the major crops barley, maize rice, and wheat, while there are gaps for the remaining crops. We do not use the hand-collected data for 1961 in our main empirical analysis, but we use them to verify that the data we collect are consistent with the FAO data discussed above. The crop-level data sources are: Mitchell (1982), Rose (1985), the Yearbook of the the League of Nations (various issues), the Statistical Yearbook of the United Nations (various issues), reports from the Economic Research Service of the US Department of Agriculture, and from the national statistical agencies in China and India. We convert all observations to common metrics (hectares and tonnes). In some cases, this involves converting quantities measured at one stage of production to equivalent quantities measured at another stage of production. Rice production is, for instance, reported at different stages of processing in the different sources. How much weight a grain of rice loses when husk and bran are removed in the milling process depends on the rice variety. We use the country-specific conversion factors in Rose (1985) to convert production of milled rice into the equivalent production of unprocessed paddy rice, the metric of production used by FAO.

D.1.3 Crop prices

We use price data to aggregate crop-level production and yields to corresponding aggregate figures, and to construct predicted aggregate gain in yields from HYVs under the assumption of fixed allocations of land and labor (i.e., $\widehat{GR}_{it}$). Crop-price data are only available from 1966 and onward in the FAO data. We use 1966-crop price data to aggregate across crops and year-crop prices as an outcome variable. From these data, we compute international average farm gate prices of each crop, which we use in our aggregation. No crop is grown in all countries, meaning that we cannot choose a single benchmark crop to calculate relative prices in order to purge differences in nominal price levels. Instead, we regress the log price of each crop on crop dummies and country dummies in a pooled sample of all crops in 1966. The country dummies capture the local price level and the exchange rate, so the coefficients on the crop dummies correspond to the average relative prices we need for the purpose of aggregation. We do not have price data for three minor crops in these groups: triticale, bambara beans, and roots and tubers not elsewhere included. We assume them to have prices similar to close substitutes. Our results are virtually unaffected by this assumption because the three crops are not widely grown.

D.1.4 Nutritional values

We use nutritional values of crops to aggregate crop yields, instead of the price data described above, in a robustness check. We measure nutritional value as the caloric content per unit of weight, obtained from the National Nutrient Database for Standard Reference 1 Release April, 2018, published by Agricultural Research Service, United States Department of Agriculture. Caloric content per unit of weight depends on how crops are processed, and the database we use for caloric content therefore has multiple entries for each crop. We make sure to pick the entry that matches the stage of processing at which FAO measure crop production, which is why our caloric content data in some cases differ from Galor and Özak (2016). Appendix Table A2 provides an overview of the caloric content of the food crops in our sample of food crops crops.

D.1.5 High-quality subsample

Many observations in the FAO data are estimates. Those produced by FAO for countries with no official statistics are flagged. The official data provided by individual countries are a mix of actual data and estimates, and estimates are not labeled. Our high-quality data set excludes FAO estimates for countries with no official statistics and official data that are 1) unvarying from year to year, 2) have less than three significant digits (we remove all zeroes trailing the last non-zero digit and count the number of digits left), and 3) are based on crops for which harvested area is less than 1000 heactares. The two first restrictions eliminate data points that are crude estimates, the last restriction removes outliers from small sample sizes in the underlying surveys.

Appendix Table A2: Caloric content of crops

Crop	Type	Kcal/g	Crop	Type	Kcal/g
Bambara beans [a]	*Pulses*	4.25	*Oats*	*Cereals*	3.89
Barley	*Cereals*	3.52	*Pigeon peas*	*Pulses*	3.43
Buckwheat	*Cereals*	3.43	*Plantain*	*Fruits*	1.22
Canary seed	*Cereals*	3.99	*Potatoes*	*Roots and tubers*	0.77
Cassava	*Roots and tubers*	1.60	*Pulses, NES* [b]	*Pulses*	3.50
Cereals, NES [b]	*Cereals*	3.50	*Quinoa*	*Cereals*	3.68
Chickpeas	*Pulses*	3.78	*Rice* [e]	*Cereals*	2.52
Cowpeas	*Pulses*	3.36	*Roots and tubser, NES* [b]	*Roots and tubers*	1.00
Dry beans	*Pulses*	3.64	*Rye*	*Cereals*	3.38
Dry peas	*Pulses*	3.52	*Sorghum*	*Cereals*	3.29
Faba beans	*Pulses*	3.41	*Soybeans*	*Pulses*	4.46
Fonio [c]	*Cereals*	3.78	*Sweetpotatoes*	*Roots and tubers*	0.86
Grain, mixed [b]	*Cereals*	3.50	*Taro*	*Roots and tubers*	1.12
Groundnut [d]	*Pulses*	4.25	*Triticale*	*Cereals*	3.36
Lentils	*Pulses*	3.58	*Vetches* [f]	*Pulses*	3.50
Lupins	*Pulses*	3.71	*Wheat*	*Cereals*	3.40
Maize	*Cereals*	3.65	*Yams*	*Roots and tubers*	1.18
Millet	*Cereals*	3.82	*Yautia* [g]	*Roots and tubers*	1.12

Notes: *Caloric content per unit of weight. Data are directly obtained from the National Nutrient Database for Standard Reference 1 Release April, 2018, published by Agricultural Research Service, United States Department of Agriculture, except in the cases noted below in which no data were available, or in which adjustments were needed to make the definitions consistent with the FAO data.* [a]*Bambara beans are related to groundnut, so we assume the same caloric content.* [b]*Residual groups are assumed to have caloric content approximately equal to the average of crops of the same type.* [c]*Fonio is a variety of millet, so we assume the same caloric content as millet.* [d]*FAO reports groundnut data with shell, whereas the nutrition data exclude the shell. We assume that shells account for 25 percent of the weight, and adjust the nutritional data accordingly.* [e]*We convert nutritional data on medium grain milled rice into caloric content of paddy rice using a conversion factor of 0.7.* [f]*Assumed to have approximately the same caloric content as other pulses (although it is primarily used for animal feed).* [g]*We use caloric content of Taro, which is botanically almost identical to Yautia.*

D.2 Country level

D.2.1 Aggregate agricultural data

The country-level agricultural outcomes, used as outcomes in Table 3 of the paper, are mainly aggregated from the the crop-level data described above. The variables are:

- *yield*: yields of food crops, defined as production per hectare of harvested land. We aggregate from the crop-level to the country-level using 1966 international prices (see above). This is our main measure of aggregate agricultural productivity (used in Panel A of Table 3).

- *food crop harvested area* (Panel B of Table 3): the total area devoted to food crops (cereals, pulses, and roots and tubers)
- *total harvested area* (Panel C of Table 3): total harvested area of all crops in the FAO data, including both food crops and cash crops. Beyond a few major cash crops, such as sugar and cotton, the data for cash crops are generally of lower quality and more frequently unbalanced than data for food crops.
- *HYV area share* (Panel D of Table 3): the estimated share of land planted with HYVs from Evenson and Gollin (2003b). These data are only available for 11 major crops, so *HYV area share* is the fraction of land planted with HYVs of these 11 crops relative to total land planted with these 11 crops (total land includes land planted with HYVs and land planted with traditional varieties). Data are only reported at a five-year frequency. We linearly interpolate between data points.
- *agricultural employment share* (Panel E of Table 3): the employment share in agriculture obtained from Wingender (2014).
- *agricultural productivity* (Panel F of Table 3): the net production value of agricultural production (in constant 2004 US$), taken from FAOSTAT ("Value of Agricultural Production"), divided by the numbers of agricultural workers. We derive the latter from the agricultural employment share and the working-age population, defined as people between 15 and 64 years.

D.2.2 Economy-wide data

Our two main country-level non-agricultural outcomes are logged *GDP per capita*, defined as real GDP per capita in constant 2011 US$, and logged population size. Both variables are obtained from the Maddison Project Database (Inklaar et al. 2018), which is more comprehensive in the pre-Green Revolution period than other sources of income data. These data are, however, supplemented with data from the Penn World Table version 9.0 (Feenstra, Inklaar, and Timmer 2015) for the following countries: Angola, Algeria, Bhutan, Fiji, and Suriname, which are missing or partially missing in the Maddison Project Database. We obtain very similar results using data only from Maddison or Penn World Table for the baseline sample period from 1961 to 2010.

The initial production shares of wheat, rice, and maize we use in the event-study estimation framework (e.g., Figure 4 in the paper), are based on production quantities from 1961-1964 using 1966 international prices for aggregation. The denominator is total production of all food crops measured in 1966 international prices. For example, the rice production share is calculated as the value of total rice production from 1961 to 1964 divided by the value of all food production over the same period. We calculate the initial production shares using data from multiple years to average out fluctuations in relative yields due to weather shocks.

In Table 6 of the paper, we study how the Green Revolution affected the following outcomes, obtained from the World Development Indicators:

- The *adult mortality rate*: the probability of dying between the ages of 15 and 60. Mortality rates are in the World Development Indicators only available by sex, so we approximate the total mortality rate by the unweighted average of the mortality rates for males and females.
- *infant mortality rate*: the number of children dying before the age of one per 1000 live births.
- *total fertility rate*: the number of children a women, in a given year, is expected to have if she was subject to the prevailing age-specific fertility rates and no risk of dying.
- The *rate of natural increase*: the crude birth rate minus the crude death rate (also known as the natural population growth rate).
- The *net migration rate*: the net inflow of people into a country. It is constructed as total population growth minus natural population growth.
- *years of schooling*, *age* $15 - 19$: the average years of schooling attained by people aged 15-19 year. Data are from Barro and Lee (2013), and they are reported at a five-year frequency. We linearly interpolate between data points to get annual data.

In Appendix Table A13 below, we study the following additional outcomes, also obtained from the World Development Indicators:

- *physicians/capita*: the number of medical doctors per capita.
- *hospital beds/capita*: the number of hospital beds per capita.
- *telephone subscriptions/capita*: the number of fixed telephone-lines subscriptions per capita.

- *electricity consumption/capita* (kWh per capita): the electricity output of power plants (including combined heat and power plants) less transmission, distribution, and transformation losses, measured in per capita terms.
- *development aid/capita*: net official development assistance received (ODA) per capita.
- *government spending/GDP*: general government final consumption expenditure as percentage of GDP.
- *manufacturing/GDP*: manufacturing value added as a percentage of GDP.

These variables are available at an annual frequency, but they are often unbalanced in the time dimension. The electricity data only start in the 1970. We fill in gaps in the data using linear interpolation.

We use annual temperature and precipitation data to control for weather shocks and climate (the "climate change" controls in Tables 3 and 4). The temperature data are from FAOSTAT, while the precipitation date are from the World Bank. To take possible non-linear effects into account, we bin the temperature/precipitation data into 10 equal sized bins and omit the fourth bin, and thereby control for climate change non-parametrically.

The geographical control variables used in Appendix Table A10 below are:

- *latitude*: absolute latitude.
- *arable land*: the percentage of a country's landmass where agriculture is possible.
- *distance coast/river*: average distance to a coast or a river from any point in a country
- *landlocked*: an indicator variable for whether a country is landlocked.
- *Initial malaria risk*: percent of population at risk of contracting malaria in 1960
- *%desert*: percent of the country classified as desert
- *%tropical*: percent of the country located in tropical climate

These variables are obtained from Galor and Özak (2016), Nunn and Puga (2012), and Gallup and Sachs (2001). Additionally, we control for trade liberalization using data from Sachs et al. (1995). The Sash and Warner (SW) indicator was originally constructed as a measure of openness to international trade, but, as argued by Buera, Monge-Naranjo, and Primiceri (2011), "the SW indicator is better interpreted as a broader, albeit stark, measure of market orientation." We also control for international trade using data from Feyrer (2009) by the "simple" instrument (applied in his long-differences specifications), based on the average of the log difference between air and sea trade distances from country i to all trading partners (countries).

E Summary statistics and baseline sample

This section reports supporting material to Section 5 in the paper, describing our data and sample. Appendix Table A2 reports summary statistics for the main variables used in the country-level analysis over the sample period 1961-2010. For example, one can read off this table that the average pre-Green Revolution wheat (production) share is 12 percent, while the corresponding average pre-Green Revolution rice share is 23 percent, and the pre-Green Revolution maize share is 17 percent.

Appendix Table A3: Summary statistics (country level, 1961-2010)

Variables	(1) N	(2) mean	(3) sd	(4) max	(5) min
$\ln \widehat{GR}$	4,500	0.154	0.139	0.582	0
$\ln \widehat{GR}$ *in* 2010	4,500	0.366	0.129	0.582	0.127
$\ln \widehat{GR} \times initial\ AES$	4,300	0.103	0.0966	0.471	0
Pre-GR wheat production share	4,500	0.106	0.205	0.846	0
Pre-GR rice production share	4,500	0.219	0.291	0.960	0
Pre-GR maize production share	4,500	0.151	0.179	0.708	0
Pre-GR Wheat+Rice+Maize production shai	4,500	0.475	0.290	0.970	0.00930
ln yields	4,500	6.769	0.605	8.358	4.606
ln harvested area, total	4,500	14.49	1.654	19.07	9.243
ln harvested area, food crops	4,500	14.08	1.768	18.79	8.493
ln agricultural employment share	4,300	-0.753	0.613	-0.0513	-3.912
ln agricultural productivity	4,150	7.552	0.965	11.20	5.483
ln GDP/capita	4,473	7.999	0.901	10.35	5.063
ln population size	4,473	8.787	2.402	14.10	-1.233
ln infant mortality rate	4,189	4.179	0.731	5.589	1.253
ln fertility rate	4,500	1.593	0.400	2.182	0.0733
ln adult mortality rate	4,450	5.596	0.421	6.756	4.215
rate of natural increase (RNI)	4,500	0.0247	0.00688	0.0413	-0.0132
net immigration rate	4,450	5.596	0.421	6.756	4.215
years of schooling, 15-20	3,850	5.206	2.515	12.46	0.171

Notes: *This table reports summary statistics for main variables used in the country-level analysis (1961-2010). The production shares are measured before the Green Revolution and vary only in the cross-section. See explanations and sources in Appendix D.*

Appendix Table A4 lists the 90 developing countries included in our baseline sample. As can be seen from the table, Angola, Bhutan, and Suriname are missing nine years of income data over this sample period (Comoros, Guinea-Bissau, Mauritania, and Swaziland are not included in specifications based on Equation 14, due to missing data on the agricultural employment share in 1961). The locations of the 90 countries are shown in the maps in Figure 3 of the main paper. Appendix Table A4 also lists which of the sub-samples in Table

6 a country belongs to (1: LDCs as of today; 2: low GR effect tertile; 3: high GR effect tertile; 4: low initial income tertile; 5: high initial income tertile).

Appendix Table A4: Countries in the baseline sample

country	*obs.*	*sample*	*group*	*country*	*obs.*	*sample*	*group*	*country*	*obs.*	*sample*	*group*
Afghanistan	50	base	1,3	*Lebanon*	50	base	2,5	*Barbados*	50	Size	na
Algeria	50	base	3,5	*Lesotho*	50	base	1,3,4	*Belize*	41	Size	na
Angola	41	base	4	*Liberia*	50	base	1,3,4	*Brunei*	41	Size	na
Argentina	50	base	2,5	*Madagascar*	50	base	1,3	*Oman*	50	Size	na
Bangladesh	50	base	3,4	*Malawi*	50	base	1,3,4	*Trinidad and Tobago*	50	Size	na
Benin	50	base	1,4	*Malaysia*	50	base	5	*Botswana*	50	Res.	na
Bolivia	50	base	2	*Mali*	50	base	1, 4	*Equatorial Guinea*	50	Res.	na
Brazil	50	base	5	*Mauritania*	50	base		*Gabon*	50	Res.	na
Burkina Faso	50	base	1, 4	*Mexico*	50	base	5	*Iraq*	50	Res.	na
Burundi	50	base	4	*Mongolia*	50	base	3	*Libya*	50	Res.	na
Butan	41	base	3,4	*Morocco*	50	base		*Saudi Arabia*	50	Res.	na
Cabo Verde	50	base		*Mozambique*	50	base	1, 4				
Cambodia	50	base	1,3,4	*Myanmar*	50	base	3,4				
Cameroon	50	base		*Namibia*	50	base	2,5				
C. African Republic	50	base	1,2,4	*Nepal*	50	base	1,3,4				
Chad	50	base	1,2,4	*Nicaragua*	50	base	3,5				
Chile	50	base	2,5	*Niger*	50	base	1,2,4				
China	50	base	4	*Nigeria*	50	base	2				
Colombia	50	base	2,5	*Pakistan*	50	base	3,4				
Comoros	50	base	1, 4	*Panama*	50	base	5				
Congo	50	base	2,5	*Paraguay*	50	base	2				
Costa Rica	50	base	5	*Peru*	50	base	2,5				
Côte d'Ivoire	50	base	5	*Philippines*	50	base	3				
Cuba	50	base	2,5	*R. of Korea*	50	base					
D.R. of the Congo	50	base	1,2	*Rwanda*	50	base	1,2,4				
Dominican R.	50	base	5	*Senegal*	50	base	1,2				
Ecuador	50	base	5	*Sierra Leone*	50	base	1,3,4				
Egypt	50	base	5	*South Africa*	50	base	5				
El Salvador	50	base	3,5	*Sri Lanka*	50	base					
Ethiopia	50	base	1, 4	*Sudan*	50	base	2				
Fiji	50	base	2,5	*Surinam*	41	base	2,4				
Gambia	50	base	1,2,4	*Swaziland*	50	base					
Ghana	50	base	2	*Syrian Republic*	50	base					
Guatemala	50	base	3,5	*Thailand*	50	base	3				
Guinea	50	base	1, 4	*Togo*	50	base	1,2				
Guinea-Bissau	50	base	4	*Tunisia*	50	base	3				
Haiti	50	base	1	*Turkey*	50	base	3,5				
Honduras	50	base	3	*Tanzania*	50	base	1, 4				
India	50	base	4	*Uganda*	50	base	1,2,4				
Indonesia	50	base	3	*Uruguay*	50	base	5				
Iran	50	base	5	*Venezuela*	50	base	2,5				
Jamaica	50	base	2,5	*Viet Nam*	50	base	3,4				
Jordan	50	base	2,5	*Yemen*	50	base	2				
Kenya	50	base	3	*Zambia*	50	base	3				
Lao People's DR	50	base	3,4	Zimbabwe	50	base	1,3				

Notes: *"Obs." refers to the number of years GDP per capita is observed in the sample 1961-2010. Base indicates whether the country is included in the baseline sample. Comoros, Guinea-Bissau, Mauritania, and Swaziland are not included in specifications based on Equation 14, due to missing agricultural employment share data in 1961. "Size" (or "Res.") indicates that the country is excluded in the base sample because of a small agricultural sector (or natural recourse dependence). "Group" refers to Table 6 of the paper: 1 LCDs; 2 Low GR; 3 high GR; 4 low initial income; 5 high initial income.*

F Additional crop-level results

This section reports supporting material to Section 6.1 in the paper, which investigates how the release of HYVs impacted crop yields. In the main paper, we report crop estimates under the assumption that the effect of HYVs on yields is identical across all treated crops (Panel A of Figure 2), and under the assumption that it is identical for crops of the same type (cereals, pulses, and roots and tubers, corresponding to Panels B-D of Figure 2). In Appendix Figure A1, we show event-study estimates for individual crops (wheat, rice, maize, barley, millet, sorghum, cassava, sweet potatoes, and yams). We do not show event-study estimates for pulses as a group rather than individually, as pulses are relatively minor crops for which our data contain insufficient observations to produce meaningful event-study graphs at the individual crop level. The event-studies reveal that the largest impacts of HYVs on yields are for wheat, rice, maize, and barley, indicating that these crops to a large extent drive the event-study estimates for all crops reported in Panel A of Figure 2.

Appendix Figure A1: Event-study estimates for each crop

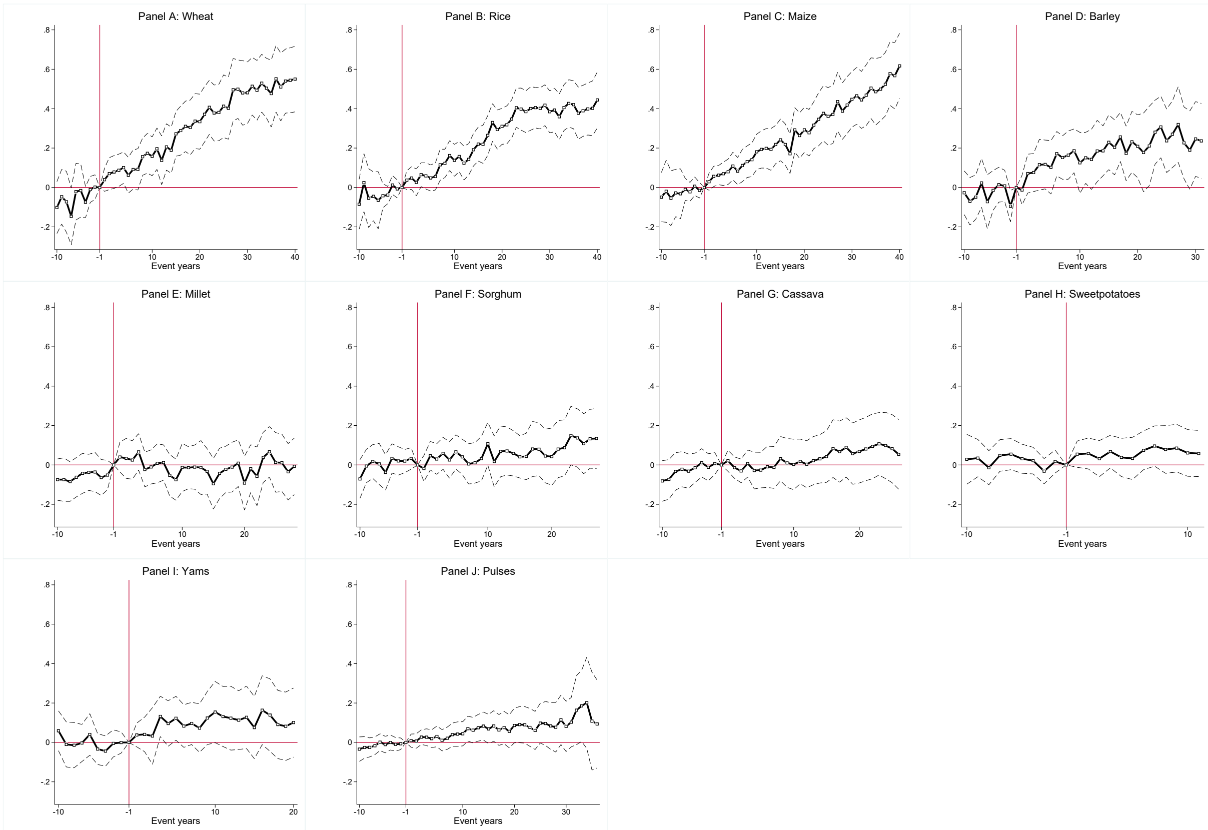

Notes: *This figure shows event-study estimates from Equation 10 for the treated crops: wheat, rice, maize, barley, millet, sorghum, cassava, sweet potatoes, yams (Panels A-I) and for pulses as a group (Panel J). The dependent variable is* ln *yields. All event-studies are estimated in a single regression based on data from the period is 1945-2010. The regression include controls for:* $\ln harvestedarea$*, country-by-year fixed effects, and crop-by-country fixed effects. The omitted comparison event-year is the year prior to the first HYV release, which varies across crops. The dashed lines are 95-percent confidence bands based on standard errors clustered at the country-crop level.*

In Appendix Figure A2, we control for pre-Green Revolution crop-specific linear pre-trends, but otherwise the specification is similar to the one used in the previous figure, Appendix Figure A1. Following Bhuller et al. (2013) and Goodman-Bacon (2018), we estimate crop-specific linear pre-trends based on yields prior to the Green Revolution (i.e., before 1965). Including linear trends based on all sample years (i.e., both pre- and post-Green Revolution) will absorb some of the actual treatment effect, particularly if the treatment effects are increasing over time, as they are in our case; see (Goodman-Bacon 2018). The event-study estimates are pretty similar with and without controls for pre-trends, suggesting that our findings are not driven by (linear) pre-Green Revolution trends.

Appendix Figure A2: Event-study estimates for each crop with controls for pre-GR trends

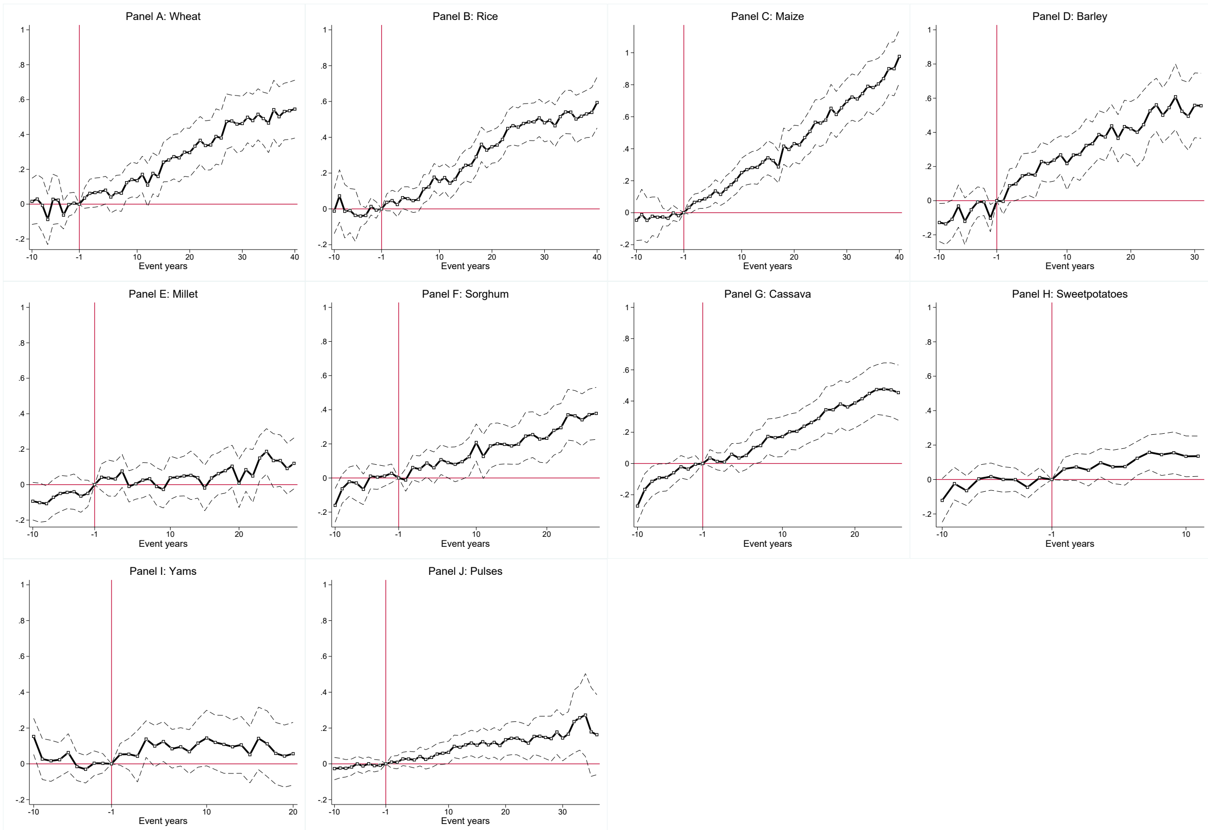

Notes: *This figure shows event-study estimates based on Equation 10 with controls for crop-specific linear pre-trends for the treated crops: wheat, rice, maize, barley, millet, sorghum, cassava, sweet potatoes, yams (Panels A-I) and for pulses as a group (Panel J). The dependent variable is* ln *yields. All event-studies are estimated in a single regression based on data from the period is 1945-2010. The regression include controls for:* $\ln harvestedarea$*, country-by-year fixed effects, and crop-by-country fixed effects. The omitted comparison event-year is the year prior to the first HYV release, which varies across crops. The dashed lines are 95-percent confidence bands based on standard errors clustered at the country-crop level.*

In Appendix Table A5, we report individual differences-in-differences estimates for each treated crop. Again, we find that wheat, rice, maize, and barley have the largest increases in yields due to HYVs. We use the individual estimates, reported in Column 1 of Appendix Table 5A, to construct the predicted yields in our country-level analysis, but we obtain similar results when using estimates based on the assumption that effects are identical across all treated crops, and when using estimates based on an assumption of identical effects across all crops belonging to the same group (cereals, pulses, roots and tubers). We do not report the individual crop estimates in the main paper to save space.

Appendix Table A5: Crop-specific estimates

	(1)	(2)	(3)	(4)	(5)
HYV x t, wheat	0.0139***	0.0148***	0.0133***	0.0136***	0.0114***
	(0.00174)	(0.00165)	(0.00189)	(0.00188)	(0.00237)
HYV x t, rice	0.0115***	0.0128***	0.0111***	0.0111***	0.00891***
	(0.00145)	(0.00160)	(0.00151)	(0.00157)	(0.00216)
HYV x t, maize	0.0141***	0.0159***	0.0142***	0.0130***	0.0116***
	(0.00176)	(0.00202)	(0.00187)	(0.00172)	(0.00242)
HYV x t, barley	0.00982***	0.00875***	0.0102***	0.0106***	0.00645*
	(0.00275)	(0.00306)	(0.00282)	(0.00291)	(0.00333)
HYV x t, millet	0.00235	0.00321	0.00220	0.00123	-0.00123
	(0.00264)	(0.00247)	(0.00265)	(0.00277)	(0.00316)
HYV x t, sorghum	0.00645**	0.00555**	0.00609**	0.00638**	0.00280
	(0.00255)	(0.00228)	(0.00254)	(0.00277)	(0.00304)
HYV x t, cassava	0.00681**	0.00630**	0.00677**	0.00680*	0.00758**
	(0.00325)	(0.00315)	(0.00326)	(0.00352)	(0.00379)
HYV x t, groundnut	0.00643***	0.00601**	0.00616**	0.00597**	0.00501
	(0.00245)	(0.00248)	(0.00246)	(0.00267)	(0.00313)
HYV x t, drybeans	0.00616***	0.00635***	0.00608***	0.00582***	0.00500**
	(0.00181)	(0.00191)	(0.00181)	(0.00195)	(0.00240)
HYV x t, lentils	-0.00128	0.000753	-0.00149	-0.000798	-0.00247
	(0.00345)	(0.00272)	(0.00345)	(0.00418)	(0.00385)
HYV x t, cowpeas	0.00535	0.0110**	0.00521	0.00366	0.00434
	(0.00359)	(0.00507)	(0.00357)	(0.00358)	(0.00390)
HYV x t, pigeonpeas	0.00538	0.00402	0.00579	0.00710	0.00360
	(0.00579)	(0.00854)	(0.00562)	(0.00592)	(0.00580)
HYV x t, chickpeas	0.00481	0.00704**	0.00461	0.00428	0.00344
	(0.00357)	(0.00357)	(0.00355)	(0.00399)	(0.00398)
HYV x t, sweetpotatoes	0.00819	0.00971	0.00794	0.00929	0.00956*
	(0.00590)	(0.00655)	(0.00586)	(0.00610)	(0.00554)
HYV x t, yams	0.00633	0.0100*	0.00612	0.00542	0.00713
	(0.00417)	(0.00556)	(0.00415)	(0.00454)	(0.00481)
HYV x t, fababeans	0.00492	0.000348	0.00474	0.00495	0.00343
	(0.00376)	(0.00359)	(0.00376)	(0.00429)	(0.00413)
ln harvested area	-0.0278***	-0.0199	-0.0333***	-0.0310***	-0.0278***
	(0.0101)	(0.0156)	(0.0105)	(0.0109)	(0.0101)
country x crop FE	Yes	Yes	Yes	Yes	Yes
crop-type linear trends	No	No	No	No	Yes
Observations	45,184	27,015	43,256	40,372	45,184
Countries	90	89	90	84	90

Notes: *This table reports crop-level estimates from Equation 11. The dependent variable is* ln *yields. The explanatory variable is an indicator variable for years after the first release of a HYV of a given crop, interacted with a linear year trend* ($HYV \times t$)*. This variable is further interacted with a set of dummies for the individual crops in our sample (e.g., wheat and rice). This table is similar to Table 2 in the paper, except that we here report estimates by individual crops (instead of aggregated, or by crop type). The sample period is 1945-2010, but unbalanced with fewer observations before 1961. Column 1 corresponds to our baseline specification. Column 2 uses only high quality data as defined in footnote 21 in the main text. Column 3 uses only FAOSTAT data (so the sample period is 1961-2010). Column 4 excludes yield data from countries with an international agricultural research center. All regressions control for* $\ln harvestedarea$*, country-by-year fixed effects, and crop-by-country fixed effects, and Column 5 additionally controls for crop-type specific linear trends (i.e., linear trends for cereals, pulses, and roots and tubers). Standard errors are clustered at the crop-by-country level.* $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

In Appendix Table A6, we report crop-level results for other outcomes than yields. In Columns 1-4, we investigate whether the increases in yields following the Green Revolution were due to decreasing harvested areas, increasing production, or both. We find that the harvested area of treated crops (i.e., crops of which HYVs were released) increases relative to that of untreated crops (i.e., crops for which HYVs became available on a later date, or never, in the case of pure control crops); see Columns 1-2. Note that we here estimate relative effects, meaning that this result only shows that the harvested area of treated crops grew relative to that of untreated crops, implying substitution. Columns 3 and 4 reveal that relative production is increasing by more than harvested area, which is equivalent to our baseline finding that relative yields of treated crops are increasing after a HYV release. In Columns 5 and 6 of Appendix Table A6, we provide suggestive evidence that the relative prices of treated crops decline following the introduction of HYVs. These results come with the important caveat that we were only able to obtain crop prices at the country-level starting in 1966, implying that we do not have pre-Green Revolution prices for the largest and most important food crops (i.e., wheat, rice, and maize). The results should consequently be interpreted with some caution as the missing pre-Green Revolution data for the three larges crops prevents us from excluding crop-specific effects.

Appendix Table A6: Additional crop-level results

	Dependent variable is (in ln)					
	Harvested area		Production		Crop prices	
	(1)	(2)	(3)	(4)	(5)	(6)
HYV x t	0.00734**		0.0196***		-0.0112***	
	(0.00311)		(0.00333)		(0.00199)	
HYV x t, cereals		0.00782**		0.0203***		-0.0108***
		(0.00314)		(0.00337)		(0.00202)
HYV x t, roots/tubers		0.0162***		0.0247***		-0.00329
		(0.00566)		(0.00637)		(0.00373)
HYV x t, pulses		0.00925**		0.0156***		-0.00699***
		(0.00426)		(0.00448)		(0.00270)
country x year FE	Yes	Yes	Yes	Yes	Yes	Yes
country x crop FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	45,412	45,412	45,330	45,330	24,316	24,316
Countries	90	90	90	90	78	78

Notes: *This table reports additional crop-level estimates based on Equation 11. The dependent variables are (in* ln*): harvested area (Columns 1-2), production (Columns 3-4), and crop prices (Columns 5-6). The explanatory variable is an indicator for the release year of an HYV in a given crop, interacted with a linear year trend (*$HYV \times t$*). This variable is further interacted with crop-type indicator (cereals, roots/tubers, pulses) in Columns (2)-(6). The sample period is 1945-2010 in columns 1-4, and 1966-2010 in columns 5-6. Standard errors are clustered at the crop-by-country level.*

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

G Agricultural Outcomes

This section reports supporting material to Sections 6.2 and 6.3 in the main paper, which investigate how the Green Revolution impacted the aggregate agricultural sector. We start by presenting further descriptive statistics for the predicted aggregate yield growth due to the Green Revolution; i.e., $\ln \widehat{GR}_{it}$, which is the main explanatory variable in our country-level analysis. Appendix Figure A3 displays this variable by country and year. Panel A depicts our baseline specification, revealing that HYVs, ceteris paribus, are predicted to have increased aggregate yield by approximately 14 (minimum) to 65 (maximum) percent from 1964 to 2010 (or 0.13-0.58 log points). Alternative specifications of the predicted yields are also depicted. Instead of crop-specific trends, we assume common crop-type (cereals, pulses, and roots/tubers) trends at the crop level in Panel B, while Panel C shows the predicted yields when we use caloric content rather than prices to aggregate crop production. All measures of predicted yields are highly correlated. For example, in 2010 $\rho_{A,B} = 0.95$ and $\rho_{A,C} = 0.96$.

As can be seen from the figure, the predicted yield gains vary across countries and years after 1965. This is particular clear from Panel D, where we have calculated the average annual growth rate in $\ln GR_{it}$ from 1964 up to the year indicated on the x-axis.

Appendix Figure A3: The *GR* variable

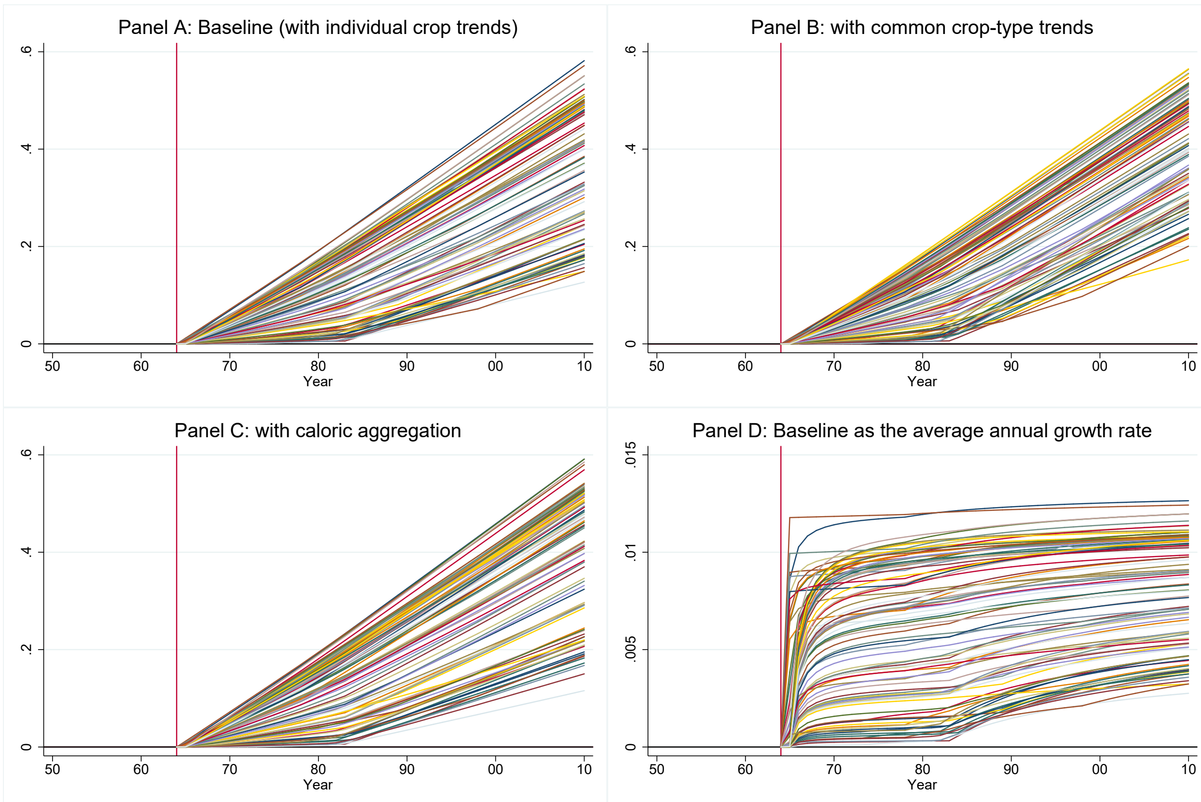

Notes: *This figure shows our predicted GR yields,* $\ln \widehat{GR}_{it}$*, by country and year. The variable can be interpreted as an approximate growth rate in yields after the Green Revolution, under the assumptions that allocations of land and labor are fixed. Panel A shows the baseline specification using individual crop estimates. Panels B and C show alternative specifications: In Panel B,* $\widehat{GR}_{it}$ *is derived assuming common yield gains across crops belonging to the same group (cereals, pulses, and roots and tubers). In Panel C, we use nutritional content rather than prices to aggregate crop-specific yield gains. Based on our baseline specification, Panel D reports the average annual growth rate from 1964 up to the year indicated on the x-axis. The vertical red line marks the last pre-Green Revolution period.*

In Appendix Figure A4, we display the interaction between $\ln \widehat{GR}_{it}$ and the initial agricultural employment share measured in 1961. We use this interaction when we estimate the effects of the Green Revolution on non-agricultural outcomes (e.g., Table 4 in the paper).

Appendix Figure A4: *ln GR x initial AES* in 2010

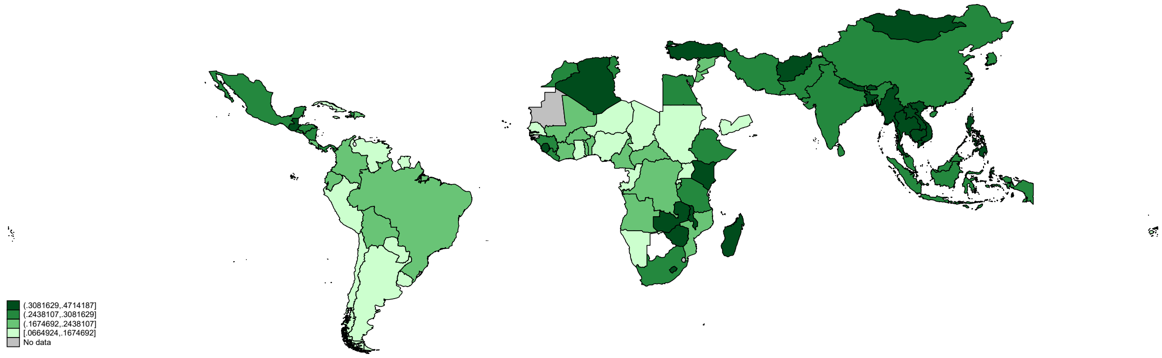

Notes: *This figure shows the value of* $\ln \widehat{GR}_{it}$ *in year 2010 interacted with the initial employment share in 1961.*

In our baseline specification, we use international crop prices from 1966 to aggregate across crops. Ideally, we would like to use pre-Green Revolution prices as 1966 prices, in principle, could be influenced by the onset of the Green Revolution. Comprehensive data on prices are not available before 1966, however, and any effects of HYVs on 1966 crop price are likely to be modest, since the total area planted with HYVs constituted less than 0.5 percent of the total area devoted to food crops (calculations based on Evenson and Gollin (2003b) and the FAO data described above). By contrast, the number was 7.4 percent in 1970. In the main text, we provide evidence that our results are not driven by any price effects by presenting the result of a robustness check in which we use nutritional content in the aggregation instead of prices. In Appendix Figure A5, we provide further evidence to this effect. Panel A displays a strong positive relationship between predicted yields based on 1966-price aggregation (our baseline explanatory variable) and predicted yields based on nutritional content aggregation. Panel B shows a similar strong positive relationship between *observed* aggregate yields based on the two different aggregation methods.

Appendix Figure A5: Price vs. caloric aggregation

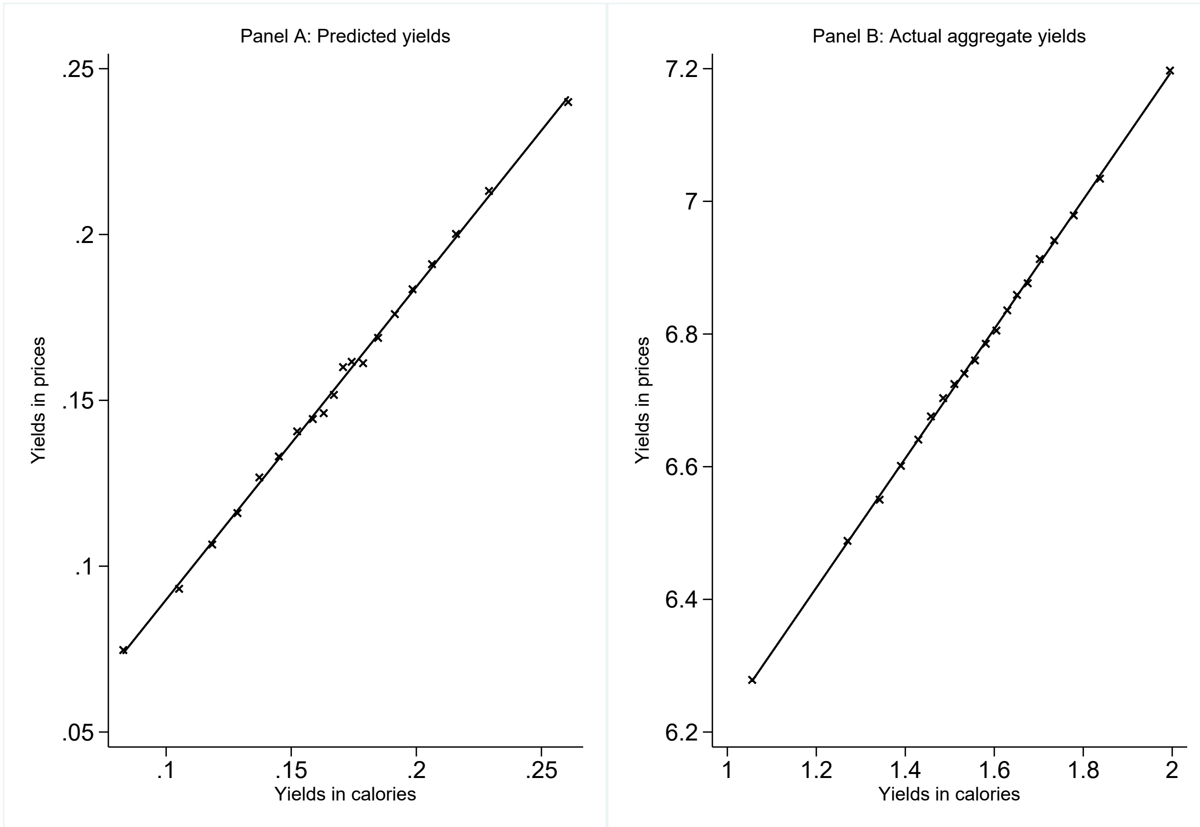

Notes: *This figure shows binned scatter plots between (predicted) yields using price and caloric aggregation, controlling for country and year fixed effects. In Panel A, we compare* $\widehat{GR}_{it}$ *derived using the two different methods of aggregation, while we in Panel B compare observed aggregate yields (endogenous yields) calculated using the two methods of aggregation.*

Appendix Figure A6 displays event-study estimates for the country-level agricultural outcomes we study in Table 3 of the main paper. The specification is similar to estimation Equation 15 in the paper, which uses sum of the production shares of wheat, rice, and maize as an initial measure of treatment intensity. The sample length is shorter than in the previous event studies because we only have data on these outcomes from 1961 to 2010. The exception is the agricultural employment share (Panel E), for which data are available for the full period 1950-2010. Consistent with the findings in Table 3, countries with higher pre-Green Revolution production shares of wheat, rice, and maize experienced increasing aggregate food-crop yields immediately after the onset of the Green Revolution in 1965 (Panel A). The event-study estimates, reported in Panels B and C, reveal that total agricultural land decreased with a lag of around 25 years (Panel C). This effect is driven by food crops (Panels B). Therefore, while HYVs increased the harvested area of treated crops relative to untreated crops, our evidence shows that total agricultural land devoted to food crops decreased. This finding is consistent with HYVs allowing a subsistence food requirement to be met with fewer hectares of crop land. In Panel D, we show that that agricultural productivity, measured as the net value of agricultural production per agricultural worker, increased more in countries with higher pre-Green Revolution production shares in wheat, rice, and maize. Since the numerator of this variable is taken directly from FAOSTAT (in contrast to $\ln yields$, which we have constructed ourselves using crop-level data), we view this finding as a robustness check for the baseline yield effects. Panel E provides evidence of structural transformation: The agricultural employment share decreased more in countries with higher initial treatment intensities during the 20th century and particularly so after the 1970s – some five years after the first HYV release. Finally, in Panel F, we provide another consistency check, which shows that countries with higher pre-Green Revolution production shares in wheat, maize, and rice had higher aggregate adoption rates of HYVs. We use HYV adoption data from Evenson and Gollin (2003a) for the period 1960-2000 (see Section D.2.1 for further details). The data are based on 11 of the 16 treated crops in our sample, and do not include any of the control crops. Moreover, missing data are in some cases coded as zeroes, leading to attenuation bias. Nevertheless, this estimated pattern suggests that we are indeed capturing effects of HYV adoption on the different outcomes.

Appendix Figure A6: Event-study estimates for the aggregate agricultural sector

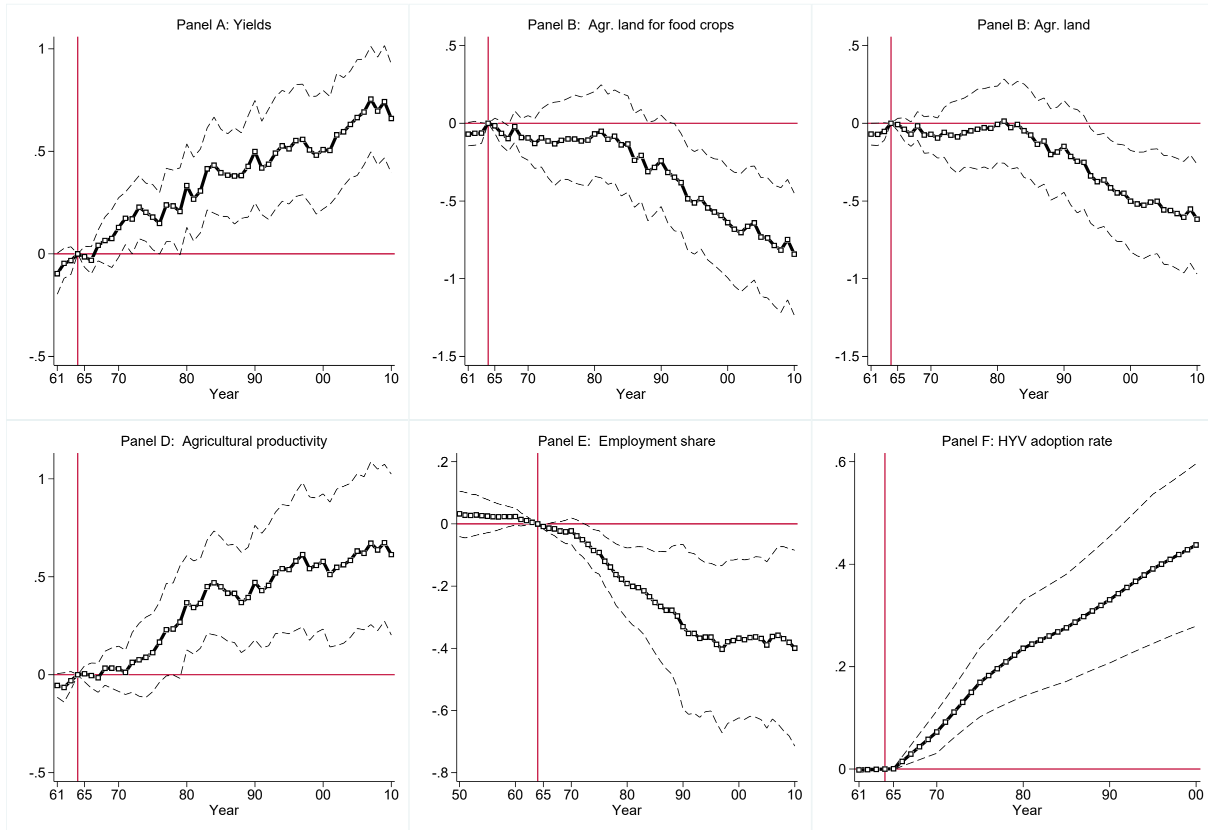

Notes: *This figure shows event-study estimates for the country-level agricultural outcomes (in* ln*): aggregate yields (Panel A); total harvested area (Panel B),; total food crops harvested area (Panel C); agricultural production per agricultural worker (Panel D); agricultural employment share (Panel E), the total HYV adoption rate (Panel F). The sample period is shorter in panels A-D and F than in Panel E because of data availability. All regressions include country and year fixed effects. The vertical red line marks the last pre-Green Revolution period, 1964, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands using standard errors clustered at the country-level.*

H Effects on development

This section contains supporting material to Section 6.4 of the paper, in which we investigate how economy-wide variables, notably GDP per capita and population size, were affected by the Green Revolution. We start out by adding controls for linear pre-trends in the main outcome variables to the baseline event-studies. In the main paper, we handled possible pre-trends using a more flexible, and more demanding, specification. Here, we take a more standard approach and control for linear pre-trends in the same fashion as Bhuller et al.

(2013): first, we estimate linear trends in outcomes in sample years before treatment. Then we add the estimated linear trend as a control for the entire sample period. This approach is also recommended by Goodman-Bacon (2018) as an alternative to directly adding a linear trend for the entire sample period to the regression, which would end up capturing a weighted average of a possible pre-trend and the relevant treatment effect.

Appendix Figure A7 reports event-study estimates for GDP per capita and population size while controlling for pre-trends using the approach explained above. The results show that, conditional on pre-trends, income growth prior to the Green Revolution was uncorrelated with production of wheat, rice, and maize (Panel A), but after the onset of the Green Revolution, countries specialized in these crops started to have higher relative income growth. The effect appears in the late 1970s and coincides with the timing of the process of structural transformation initiated by the Green Revolution, according to the event-studies for the agricultural employment share reported above.

The baseline event-study estimates for population, reported in Panel B of Figure 4 in the main text, provide evidence that pre-Green Revolution population pressure was correlated with production of wheat, rice, and maize. Such a pre-Green Revolution population trend would bias our estimate of the Green Revolution on population towards zero. Indeed, in Panel B of Appendix Figure A6, we observe that the post estimates increase in magnitude when controlling for pre-Green Revolution linear trends, while the positive pre-trend disappears.

Appendix Figure A7: Event-study estimates for income and population controlling for pre-GR trends

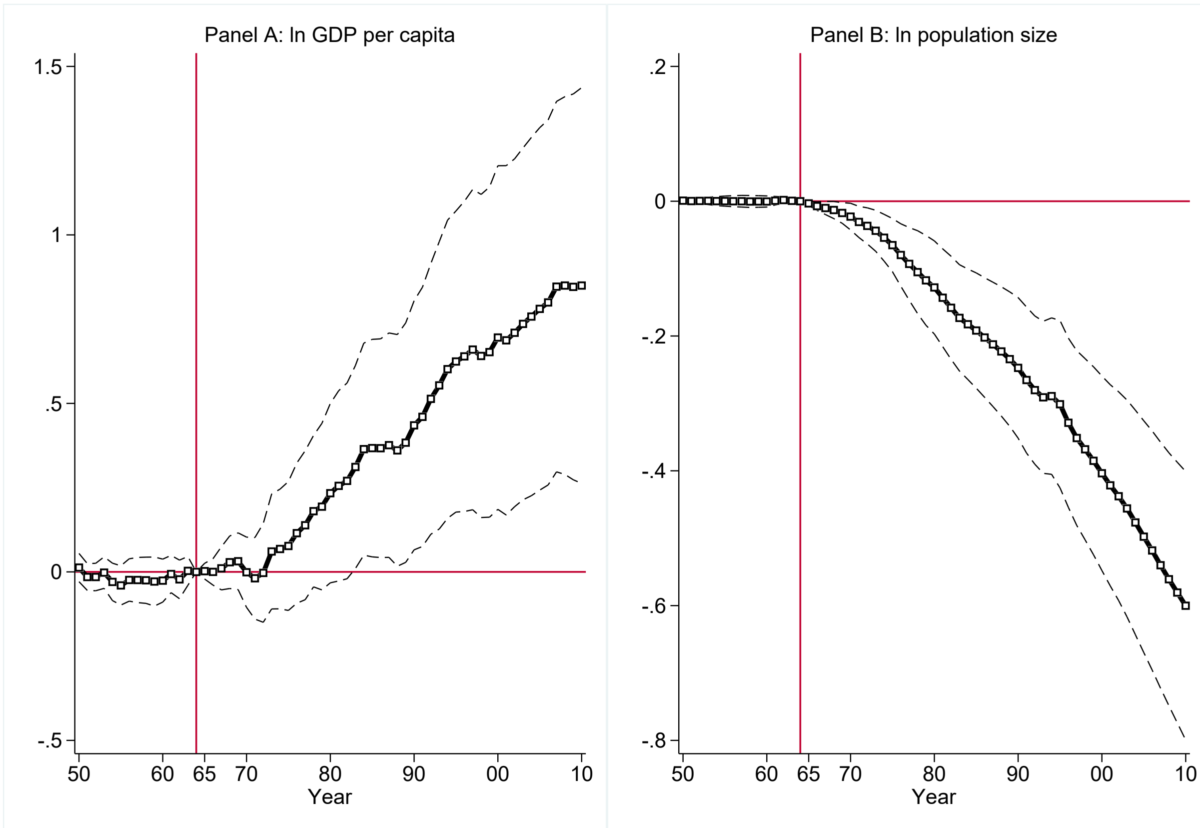

Notes: *This figure shows event-study estimates based on Equation 15 in the paper with controls for pre-Green Revolution linear trends. The dependent variables are:* $\ln GDP$ *per capita (Panel A) and* $\ln$ *population (Panel B). Both regressions include country and year fixed effects as well pre-Green Revolution linear trends. We use balanced samples in the period 1950-2010. The samples in Panel A and B have 85 and 87 countries, respectively. Similar results are obtained including all 90 countries in unbalanced samples. The vertical red lines mark 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

So far, the treatment intensity in all our country level event-studies has been the *sum* of the pre-Green Revolution production shares of wheat, rice, and maize. In appendix figures A8 and A9, we include the production shares of the three crops as separate regressors in order to construct event-study graphs for each of them. The results for income are presented in Appendix Figure 8A. The pre-Green Revolution trends are relatively parallel in all three production share, and the post-Green Revolution estimates are also positive, albeit effects are largest for wheat and rice. Therefore, if anything, this exercise informs us that countries with high pre-Green Revolution wheat and rice production shares benefited the most in terms income increases. Beyond that finding, there is not much evidence of treatment heterogeneity, and we find the same qualitative pattern for all three crops.

Appendix Figure A8: Event-study estimate for income by the crop shares: wheat, rice, and maize

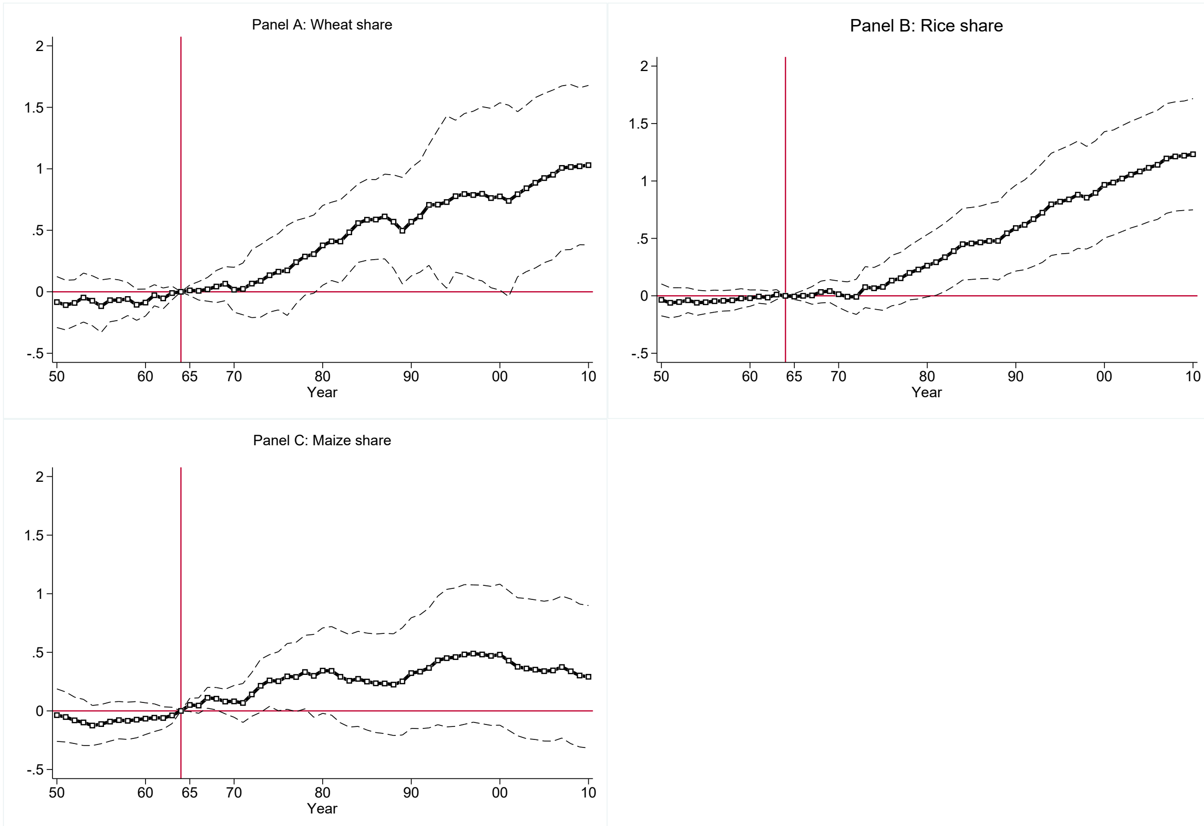

Notes: *This figure shows event-study estimates from a modified version of Equation 15 in the paper. The dependent variable is* $\ln GDP$ *per capita. Instead of summing the pre-Green Revolution production shares in wheat, rice, and maize, we here add them individually in the regression. The regression includes country and year fixed effects. We use a balanced sample of 85 countries in the period 1950-2010. Similar results are obtained including all 90 countries in unbalanced samples. The vertical red lines mark 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

The results for population, reported in Appendix Figure A9, show that the pre-Green Revolution population pressure is mainly driven by countries with high pre-Green Revolution maize production shares. By contrast, population declines are most pronounced in countries initially more dependent on rice production. As with income, we find the same qualitative pattern across all three crops.

Appendix Figure A9: Event-study estimate for population by the crop shares: wheat, rice, and maize

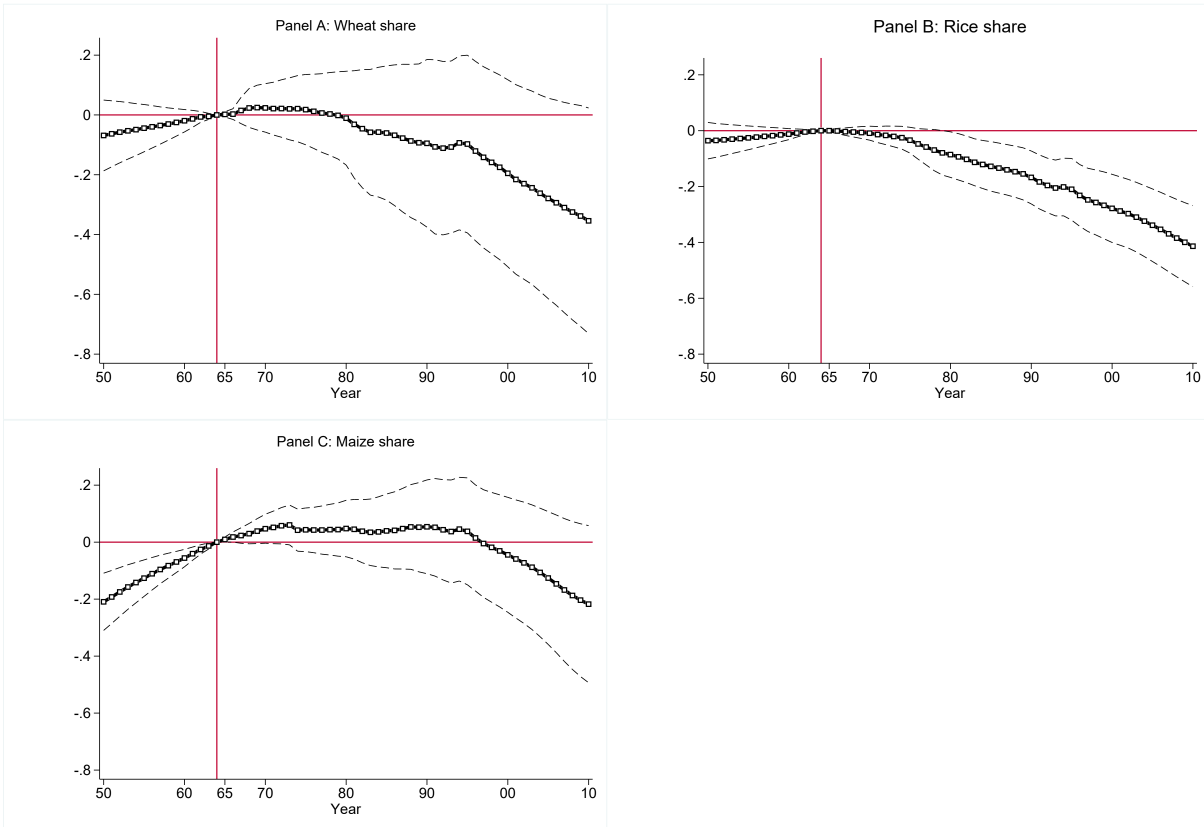

Notes: *This figure shows event-study estimates from a modified version of Equation 15. The dependent variable is* $\ln$ *population. Instead of summing the pre-Green Revolution production shares in wheat, rice, and maize, we here add them individually in the regression. The regression includes country and year fixed effects. We use a balanced sample of 87 countries in the period 1950-2010. Similar results are obtained including all 90 countries in unbalanced samples. The vertical red lines mark 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

In the baseline country event-studies, we use initial production shares as a measure of treatment intensity, whereas our baseline country regressions use HYVs contribution to yield growth ($\ln \widehat{GR}$) *interacted* with the initial agricultural employment share (L_{Ai0}/L_{i0}) as the regressor to take into account that economies more dependent on agriculture have more to gain from the Green Revolution. Since $\ln \widehat{GR} \times L_{Ai0}/L_{i0}$ is *not* simply the production shares interacted with the employment share, but rather a weighted interaction, we opted for only using the production shares in the baseline event studies. Nevertheless, in Appendix Figure A10, we show that the interaction between the sum of productions shares in wheat, maize, and rice and initial agricultural employment share is unrelated to income (and population) changes prior to the Green Revolution, and that countries with larger pre-Green Revolution production shares and more dependent on agricultural experienced larger income gains (population decreases) after 1965.

Appendix Figure A10: Event-study estimates for interaction between initial AES and production shares

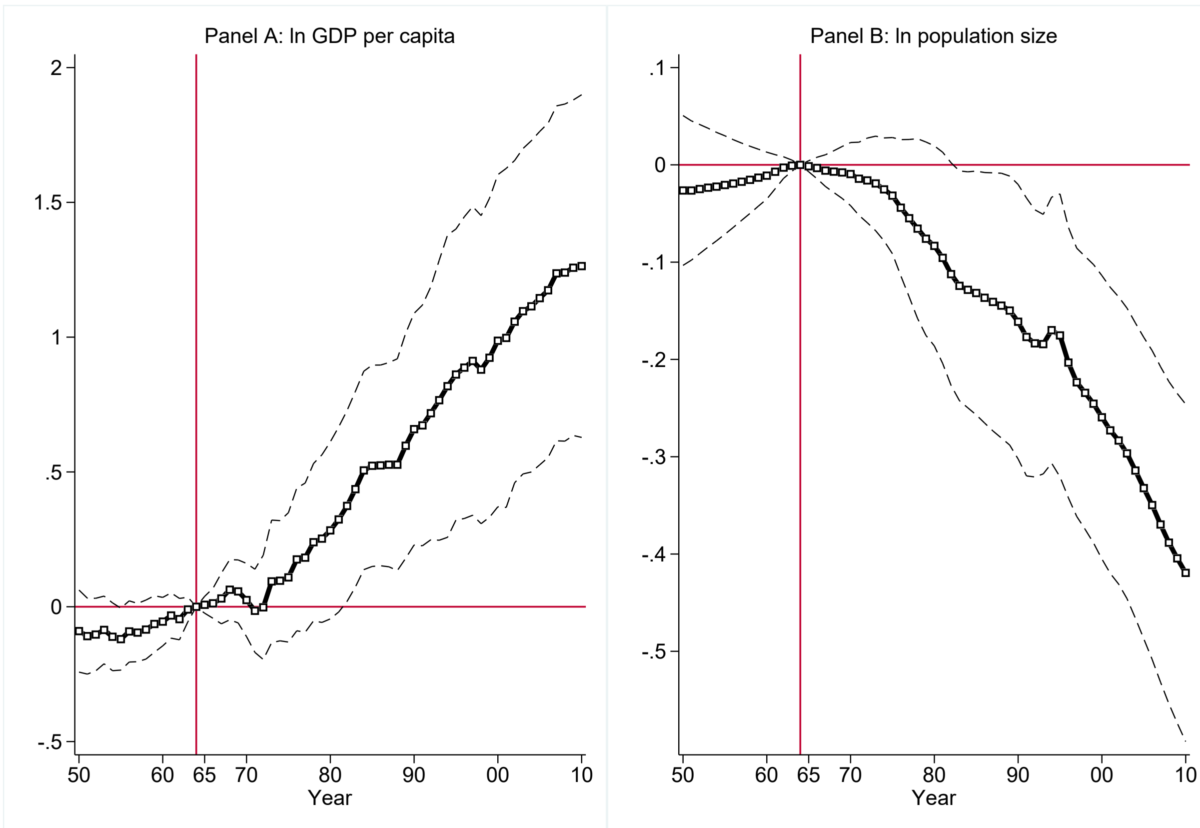

Notes: *This figure shows our country-level event-study estimates when interacting the initial production shares with initial agricultural employment shares. The dependent variables are:* $\ln GDP$ *per capita (Panel A) and* $\ln$ *population (Panel B). The explanatory variable (interacted with year fixed effects) is measured as the sum of the initial production shares in wheat, rice, and maize interacted with initial agricultural employment shares. Both regressions control for country and year fixed effects. We use a balanced sample of 85 countries in the period 1950-2010. Similar results are obtained including all 90 countries in unbalanced samples. The vertical red lines mark 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

In Appendix Figure A11, we report marginal effects of the Green Revolution on income (Panel A), population (Panel B), and agricultural productivity (Panels C and D), from a fully interacted model, including both the baseline interaction term ($\ln \widehat{GR}_{it} \times L_{Ai0}/L_{i0}$) as well as the direct effect ($\ln \widehat{GR}_{it}$). The vertical lines in the four panels indicate the median value of initial agricultural employment in the sample. As we should expect, the effects of

HYVs on income and population are larger in countries with large agricultural sectors before the Green Revolution. The effects on yields and agricultural labor productivity appear to be at most weakly related to the size of the agricultural sector, which one also what we expect to see.

Appendix Figure A11: Marginal effects of the Green Revolution

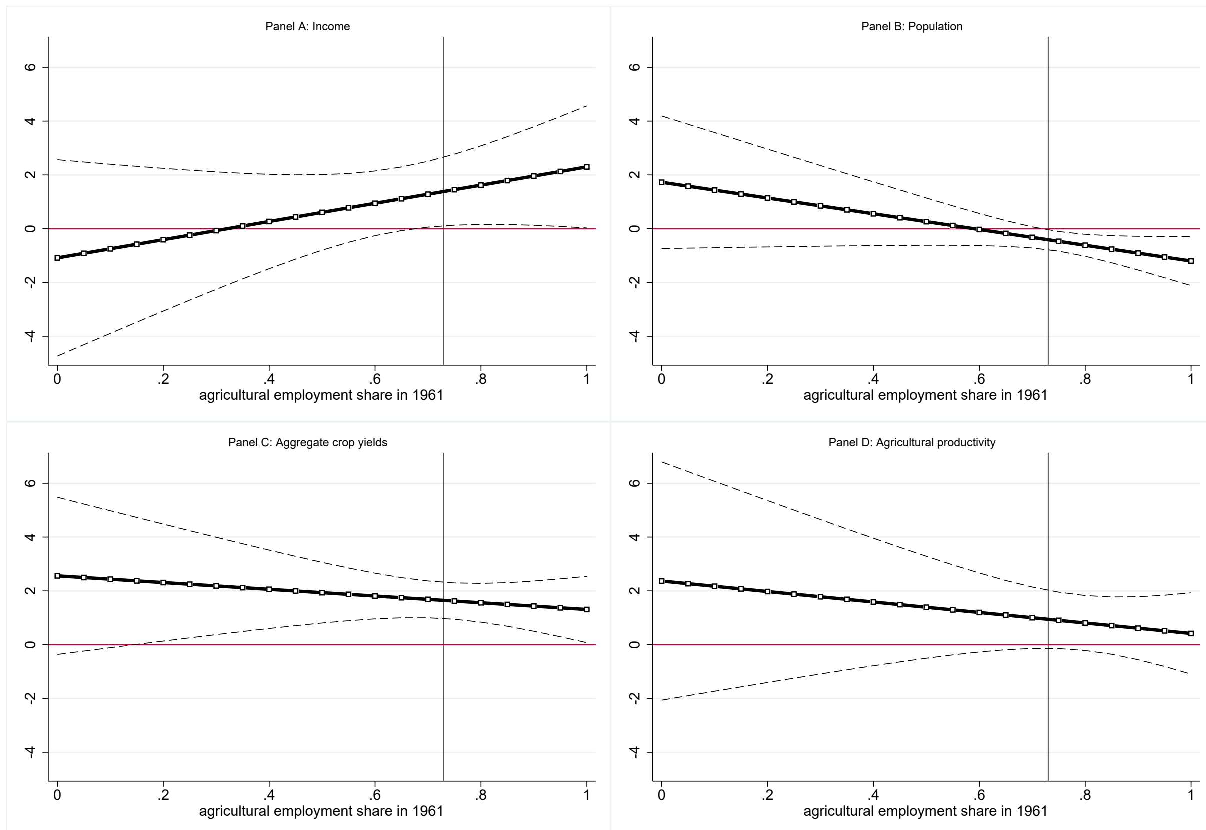

Notes: *This figure displays marginal effects of* $\ln \widehat{GR}_{it}$ *for different values of the agricultural employment share in 1961 for a fully interacted long-differences model (1961 and 2010). The dependent variables are:* $\ln$ *GDP per capita (Panel A),* $\ln$ *population (Panel B),* $\ln$ *yields (Panel C), and* $\ln$ *agricultural productivity (Panel D). The vertical line indicates the median initial agricultural employment share. The samples include only 83 countries because of data availability in 1961.*

In Appendix Table A7, we report results from falsification tests in which we ask whether the estimated yield gain affects outcome variables *before* the Green Revolution. As with the event studies in Figure 4 and Appendix figures A6-A10, the purpose is to test for possible pre-trends contaminating our results. But whereas the event-studies provide a precise picture of pre-trends and post-effects year to year, they are only based on production shares of wheat, rice, and maize. By contrast, the falsification tests in Appendix Table A7 use the entire GR prediction (interacted with the employment share), including other crops and the crop-specific yield gains we have estimated. Goldsmith-Pinkham, Sorkin, and Swift (2018) recommend such falsification tests in which one uses the "instrument" directly as a supplement to event-studies. The falsification test is less precise in terms of timing, however. To be concrete, we partition the post-Green Revolution period into three sub-periods: 1970-1980, 1965-1975, and 1960-1970, and ask if outcome changes from 1950 to 1960 are predictive of Green Revolution-driven growth in crop yields multiplied with the initial size of agricultural sector. Notice, it is possible to make this test for as many post-Green Revolution sub-periods as one wants, because $\ln \widehat{GR}$ varies across countries and years after 1965. Columns 1-4 report the findings for income, while Columns 5-8 report the findings for population. In all specifications, $\ln \widehat{GR}$ is interacted with the initial employment share. The results show that all falsification estimates are statistically insignificant. The income estimates are substantially smaller in magnitude compared to the baseline effect of 2.8, reported in Column 2 of Table 4 in the paper, while the population estimates have the opposite sign compared to the baseline effect.

Appendix Table A7: Falsification tests

	Dependent variable:					
	ln GDP/capita, 1950-1960			*ln population size, 1950-1960*		
	Post-GR sub-period					
	1970-1980	1965-1975	1960-1970	1970-1980	1965-1975	1960-1970
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln \widehat{GR} \times initial\ AES$	0.234	0.253	0.615	0.204	0.195	0.229
	(0.525)	(0.526)	(0.756)	(0.221)	(0.221)	(0.307)
Controls (x year FE):						
Initial AES	Yes	Yes	Yes	Yes	Yes	Yes
Countries	81	81	81	81	81	81
Observations	891	891	810	891	891	810

Notes: *This table reports estimates from falsification checks, which ask if* $\ln \widehat{GR}$ *interacted with the initial agricultural employment share can explain variation in outcomes prior to the Green Revolution. The dependent variable (income or population) is in all specification measured during the pre-Green Revolution period 1950-1960.* $\ln \widehat{GR}$ *x initial agricultural employment share is measured in different post-Green Revolution sub-periods indicated in the top rows (e.g., Column 1 from 1970 to 1980). This type of test is possible because the* $\ln \widehat{GR}$ *variable varies across countries and years after 1965 (see also Appendix Figure A3). All specifications include country and year fixed effect as well as the initial agricultural employment share interacted with year fixed effects. Standard errors are based on a two-step wild cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Appendix Table A8 reports estimates where we control for pre-Green Revolution linear trends. Again, we follow Bhuller et al. (2013) and Goodman-Bacon (2018) and estimate linear trends based on the pre-Green Revolution period from 1950 to 1964, and eliminate those estimated country-specific trends from the outcomes throughout the entire sample period. For this reason, Columns 1 and 4 first contain the baseline estimates when using the extended period 1950-2010. We see that adding 11 more pre-Green Revolution years does not make any significant difference compared to baseline interaction model. Controlling for pre-Green Revolution linear trends reduces the income coefficient, but it remains positive and statistically significant (Column 2). The next column shows a similar finding when restricting the sample period to 1961-2010 (Column 3). In Columns 5 and 6, we see that the estimates for population almost doubles in numerical magnitude when controlling for pre-Green Revolution linear trends, indicating that baseline population estimate is best interpreted as a lower bound.

Appendix Table A8: Controlling for pre-GR trends

	Dependent variable:					
	ln GDP per capita			*ln population size*		
	Sample period					
	1950-2010		1961-2010	1950-2010		1961-2010
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln \widehat{GR} \times initial\ AES$	2.735***	2.135*	2.252**	-0.523**	-1.051**	-1.178***
	(0.867)	(1.130)	(1.130)	(0.255)	(0.399)	(0.399)
Controls (x year FE):						
initial AES	Yes	Yes	Yes	Yes	Yes	Yes
pre-GR linear trends	No	Yes	Yes	No	Yes	Yes
Countries	86	86	86	86	86	86
Observations	5,166	5,166	4,273	5,166	5,166	4,273

Notes: *This table reports estimates for income and population, controlling for pre-Green Revolution linear trends. Columns 1 and 5 report baseline estimates, using the full period 1950-2010, since the pre-Green Revolution linear trends are estimated for the pre-Green Revolution period 1950-1964. Columns 2 and 5 report estimates controlling for pre-Green Revolution linear trends using the sample period 1950-2010, while Columns 3 and 6 report estimates using the baseline sample period 1961-2010, although the pre-Green Revolution linear trends remain based on the period 1950-1964. All regressions include country, year fixed effects, and initial agricultural employment shares interacted with year fixed effects. Standard errors are based on a two-step wild cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

In Columns 1-5 of Appendix Table A9, we demonstrate how our results vary when excluding different HYVs in the construction of $\ln \widehat{GR}$. Specifically, we exclude the following crops one at the time: wheat (Column 1), rice (Column 2), maize (Column 3), roots and tubers (Column 4), and pulses (Column 5). We are thereby moving the crops in question from the treatment measure to the control set. The income estimates, reported in Panel A, are generally robust, although excluding rice makes it insignificant. This finding does not mean that HYVs of other crops are unimportant for income, rather, it reflects that HYVs of rice now is an omitted variable. Including rice separately results in both the coefficient on $\ln \widehat{GR}$ and on rice being statistically significant (not shown in the table). The insignificance in Column 2 combined with the significance in the remaining columns is, therefore, a way of illustrating that rice is quantitatively more important for our results than the other crops. We find a similar pattern for population (Panel B). In Column 6, we use all crops in the FAO universe when constructing both predicted and actual yields, while we in Column 7 only add the most important cash crops (sugar, coffee, cotton, and tobacco). In both cases, $\ln \widehat{GR}$ becomes smaller as we now include more untreated crops, which is why the income and population estimates increase in magnitude.

Appendix Table A9: Robustness to GR variable specification

	omitted crop					including all FAO crops	including important cash crops
	wheat	rice	maize	roots	pulses	crops	cash crops
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: ln GDP per capita						
$\ln\widehat{GR} \times initial\ AES$	2.248**	0.996	2.962***	1.980**	2.470***	3.151***	3.141***
	(0.868)	(0.865)	(0.870)	(0.902)	(0.765)	(1.048)	(1.090)
	Panel B: ln population size						
$\ln\widehat{GR} \times initial\ AES$	-0.761***	0.0955	-0.758***	-0.563**	-0.429*	-0.937***	-0.684**
	(0.261)	(0.274)	(0.245)	(0.272)	(0.240)	(0.305)	(0.317)
Controls (x year FE):							
initial AES	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Countries	86	86	86	86	86	86	86
Observations	4,273	4,273	4,273	4,273	4,273	4,273	4,273

Notes: *Columns 1-5 report robustness results to the crops included in the GR variable by step-wise omitting different crops (or crop groups) in the aggregation. The top row indicate the crop being omitted. For example, Columns 1 excludes wheat and report the effects income and population. Columns 6 and 7 include additional crops in the GR variable. All specifications include country and year fixed effects. Standard errors are based on a two-step wild cluster restricted bootstrap procedure, which takes into account that* $\ln\widehat{GR}$ *is generated.*
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

We add additional control variables to our baseline specification in Appendix Table A10. As evident from Panel A (Columns 1 and 2), the income estimates are robust to controlling for logged absolute latitude, percent land in the desert, and percent land in the tropics, the initial risk of contracting malaria, and logged average distance to coast (all interacted with year fixed effects), free-market reforms as measured in Sachs et al. (1995) and Buera, Monge-Naranjo, and Primiceri (2011) (Column 4), and an instrument for trade, which is using differences in sea and air distance across countries and their trade partners, developed in Feyrer (2009), $\times$ year fixed effects (Column 5). The population estimates, reported in Panel B, are negative and of the same magnitude as in the baseline interaction model, but not when controlling for free market reforms. This main is that including the free-market reform variable restricts the sample period to 1961-1993, and we know from our baseline results that the demographic response comes with a delay such that the effect on population become negative and significant after that (see Appendix Table A11).

Appendix Table A10: Robustness to additional control variables

	Geography x year FE	All geographic variables x year FE	Initial outcome levels x year FE	Openess	Trade x year FE
	(1)	(2)	(3)	(4)	(5)
	Panel A: *ln GDP per capita*				
$\ln \widehat{GR} \times initial\ AES$	2.465**	2.157**	2.232***	2.648***	2.713**
	(0.960)	(0.971)	(0.726)	(0.970)	(1.019)
	Panel B: *ln population size*				
$\ln \widehat{GR} \times initial\ AES$	-0.300	-0.300	-0.546**	-0.00887	-0.815*
	(0.271)	(0.268)	(0.260)	(0.303)	(0.426)
	Panel C: *ln population size* w. pre-GR trends				
$\ln \widehat{GR} \times initial\ AES$	-0.642	-0.555	-0.795**	-0.455	-1.623***
	(0.432)	(0.482)	(0.325)	(0.314)	(0.478)
Controls (x year FE)					
Initial AES	Yes	Yes	Yes	Yes	Yes
ln abs latitude	Yes	Yes	No	No	No
% desert	Yes	Yes	No	No	No
% tropical	Yes	Yes	No	No	No
initial malaria risk	No	Yes	No	No	No
ln distance coast	No	Yes	No	No	No
trade	No	No	No	No	Yes
outcome variable in 1961	No	No	Yes	No	No
Observations	4,273	4,273	4,150	2,174	2,891
Countries	86	86	83	58	70

Notes: *This table reports robustness to the baseline interaction income and population estimates by including additional control variables. Column 1 controls for logged absolute latitude, percent land in desert, percent land in tropical climate (all interacted with year fixed effects). Column 2 add the initial risk of contracting malaria and logged average distance to coast (all interacted with year fixed effects) to the control set. Control variables are obtained from Galor and Özak (2016), Nunn and Puga (2012), and Gallup and Sachs (2001). Column 3 includes the outcome variable measured in 1961 interacted with year fixed effects. Column 4 includes a control for market orientation from Sachs et al. (1995). Column 5 includes the simple IV for trade (Feyrer 2009) interacted with year fixed effects. All specifications control for country, year fixed effects, and the initial agricultural employment share interacted with year fixed effects. In Panel A, the outcome variable is* ln *GDP per capita. In Panel B, the outcome variable is* ln *population size. In Panel C, the outcome variable is* ln *population size controlling for pre-Green Revolution linear trends. The sample period is 1961-2010 in all specifications, expect for Column 4 which is 1961-1993 because of data availability. Standard error are based on a two-step wild-cluster bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Our panel data do not lend themselves to graphical illustration. To inspect our our regressions visually, we plot an equivalent long-difference specification for the entire period 1961 to 2010 in Appendix Figure A12. The Figure shows that our baseline findings are not driven by outliers. In Appendix Table A11, long-difference estimates are reported with the same start date (i.e., 1961), but for different end dates (i.e., 1970, 1980,.., 2010). As can be seen from Panel A, the effect on income becomes statistically significant in the 1970s and is increasing over time. Long-difference population estimates, reported in Panel B, reveal that the effect on population only becomes negative and statistically significant starting in the 2000s, implying a delayed impact compared to income.

Appendix Figure A12: Long-difference plots for GDP per capita and population

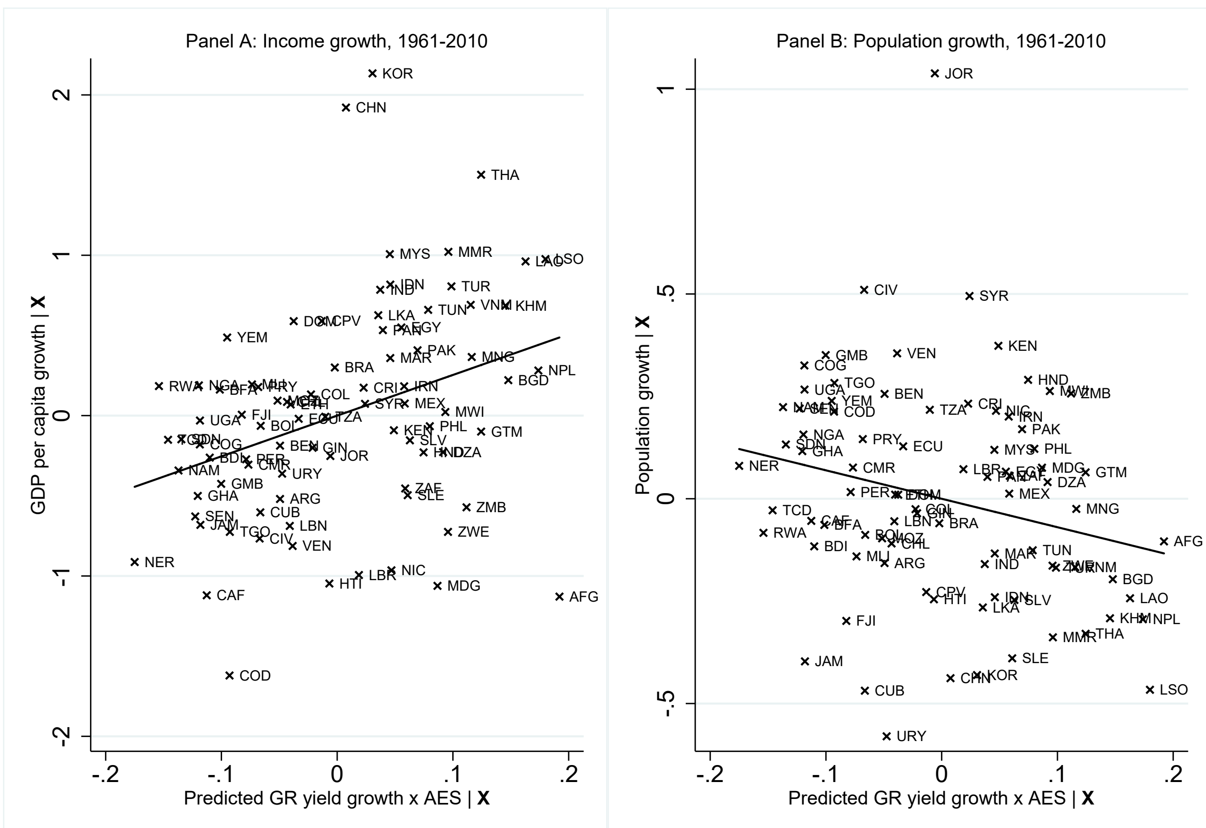

Notes: *This figure shows long-differences plots for* $\ln GDP$ *per capita (Panel A) and* $\ln$ *population (Panel A). The long difference is based on 1961 and 2010. The y-axis represents the approximate growth rates in income and population conditional on the initial agricultural employment share and the regression constant, and the x-axis represents the approximate growth rate in predicted GR yields interacted with the initial agricultural employment share. X = regression constant and the initial agricultural employment share. The sample consists of 83 countries, fewer than in our baseline regressions, because of missing data for 1961.*

Appendix Table A11: Long differences w. different end dates

	Start date and end date				
	1961 and 1970	1961 and 1980	1961 and 1990	1961 and 2000	1961 and 2010
	(1)	(2)	(3)	(4)	(5)
	Panel A: *ln GDP per capita*				
$\ln \widehat{GR} \times initial\ AES$	1.033	1.816**	2.415***	2.506***	2.538***
	(1.132)	(0.867)	(0.772)	(0.935)	(0.826)
	Panel B: *ln population size*				
$\ln \widehat{GR} \times initial\ AES$	0.0640	-0.261	-0.361	-0.446*	-0.695**
	(0.356)	(0.314)	(0.331)	(0.265)	(0.276)
Controls (x year FE)					
Initial AES	Yes	Yes	Yes	Yes	Yes
Countries	83	83	83	83	83
Observations	166	166	166	166	166

Notes: *This table reports estimates from estimating long-differences specifications akin to Figure 5, but with different end dates/years. The start date is 1961 in all specifications. The end date is indicated in the top row. For example, Column 1 uses two years of data: start date 1961 in and the end date in 1970. In Panel A, the outcome is ln GDP per capita. In Panel B, the outcome is ln population size. Country fixed effects are absorbed by differencing and year fixed effects are capture by the (unreported) regression constant. We also control for the initial agricultural employment share (AES). Standard errors are based on a two-step wild-cluster bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

In the baseline sample, we exclude countries with less than 10,000 ha of land devoted to food crops, due to concerns about data quality, and countries highly dependent on sub-soil resource extraction (diamonds and oil). In Appendix Table A12, we include these countries in the sample, with the exception of countries with insufficient agricultural data to calculate $\ln \widehat{GR}_{it}$. In practice, this corresponds to excluding countries with less than 2,000 ha of crop land.[1] The table shows that the effects on the outcomes: yields, the agricultural employment share, GDP per capita, population size, the adult mortality rate, the infant mortality rate, and the total fertility rate, are generally robust to including these additional countries. A partial exception is the income estimate, which declines from 2.8 in the baseline interaction model to 1.5 with a p-value of 17 percent when the sub-soil resource economies are

[1] For comparison, 2,000 ha is approximately half the land area of the city of Oxford and one-fourth the land area of the city of Copenhagen. Countries with less cropland than this are in some sense not producing crop output at any consequential scale.

included in Panel B. As evident from the long-difference plots in Appendix Figure A13, the lower significance of the income estimates is caused by including Botswana and Equatorial Guinea. Both have low predicted yield growth due to the Green Revolution and exceptionally high income growth because discoveries of subsoil resources within the sample period. In Botswana, diamond production started in 1971 and mining became the most important sector of the economy in less than a decade. Equatorial Guinea is even more extreme. Oil was discovered in 1996, and the subsequent exploitation increased GDP per capita in Equatorial Guinea by 16 percent per year on average in the period 1966-2010, according to our data. Including such extreme cases in our baseline sample would clearly bias our results. Instead of excluding them, we can also control for natural resources. If we, for example, control for diamond extraction (averaged 1958-2000 and taken from Nunn and Puga (2012)) per capita interacted with year fixed effects, the annual panel estimate increase to 1.9 and becomes statistically significant at the five percent level. In line with the previous empirical literature on growth (e.g. Mankiw, Romer, and Weil 1992), we prefer excluding the resource countries rather than adding control variables for resource extraction, as actual natural resource extraction is potentially an endogenous control.

Appendix Table A12: Robustness to sample

	Dependent variable (in *ln*):						
	aggregate crop yields	*agricultural employment share*	*GDP per capita*	*population size*	*adult mortality rate*	*infant mortality rate*	*total fertility rate*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A: including countries with 2000-10000 HA agricultural land						
$\ln \widehat{GR}$	1.931***	-0.754*					
	(0.452)	(0.415)					
$\ln \widehat{GR} \times initial\ AES$			2.683***	-0.477*	-1.067*	-2.306***	-2.249***
			(0.849)	(0.263)	(0.543)	(0.579)	(0.452)
Observations	4,700	4,448	4,414	4,414	4,450	4,216	4,450
Countries	94	89	89	89	89	89	89
	Panel B: including resources countries						
$\ln \widehat{GR}$	1.752***	-0.965***					
	(0.395)	(0.350)					
$\ln \widehat{GR} \times initial\ AES$			1.485	-0.594**	-1.107**	-2.001***	-2.045***
			(1.041)	(0.250)	(0.557)	(0.566)	(0.450)
Observations	4,800	4,600	4,573	4,573	4,600	4,323	4,600
Countries	96	92	92	92	92	92	92
	Panel C: Including both resources and 2000-10000 HA countries						
$\ln \widehat{GR}$	1.879***	-0.803*					
	(0.428)	(0.405)					
$\ln \widehat{GR} \times initial\ AES$			1.503	-0.438*	-1.200**	-2.333***	-2.080***
			(1.025)	(0.259)	(0.545)	(0.574)	(0.441)
Controls (x year FE)							
Initial AES	No	No	Yes	Yes	Yes	Yes	Yes
Observations	5,100	4,848	4,805	4,805	4,850	4,545	4,850
Countries	102	97	97	97	97	97	97

Notes: *This table reports robustness to including in the sample countries 2,000-10,000 ha of food crop land (Panels A and C) and countries dependent on sub-soil resources (Panels B and C). The outcome variables are indicates in the top row. The sample period is 1961-2010. All specifications include country fixed effects and year fixed effects. The initial agricultural employment share interacted with year fixed effects are included in the interaction models (Columns 3-6). Standard errors are based on a two-step wild cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Appendix Figure A13: Long-difference plot for income, 1970-2010

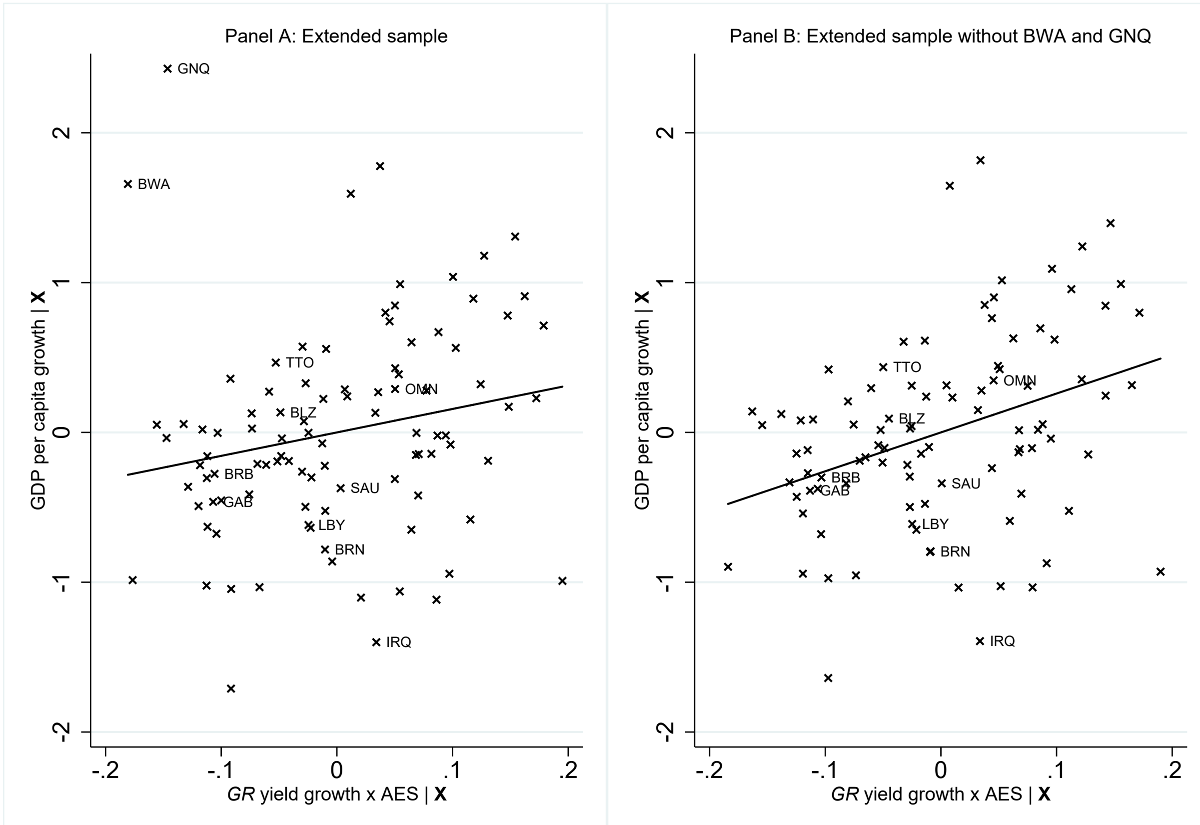

Notes: *This figure shows a long-differences plot for* $\ln GDP$ *per capita (income). We use the two years 1970 and 2010 (in order to maximize the number of countries included in the long-difference sample). The y-axis provides the approximate growth rates in income conditional on the initial agricultural employment share (AES) and the regression constant, and x-axis provides the approximate growth rate in predicted yields due to the Green Revolution times initial AES. X = regression constant and initial AES. In Panel A, all minor agricultural countries and countries highly dependent on sub-soil resource extraction are added to the baseline sample (N=97). The estimated coefficient is equal to 1.56 (p-value=0.106). In Panel B, the two outliers Botswana (BWA) and Equatorial Guinea (GNQ) are excluded and the estimated coefficient increases to 2.59 (p-value=0.001)*

In Appendix Figure 14A, we present event-study estimates for the additional non-agricultural outcomes, reported in Table 6 of the paper. While data availability restricts us to a short pre-Green Revolution period (1961-1965) for these outcomes, and prevents a rigorous pre-trend analysis, the post-Green Revolution estimates reveal that adult mortality (Panel A) started to decline later than infant mortality (Panel B); the late 1980s vs. 1970s. We also see that while fertility is declining throughout the post period (Panel C), years of schooling, age 15-20, only started to increase some 20 years later (Panel F). Consistent with the population results, the rate of natural increase is decreasing throughout the post-Green Revolution period (Panel D). There are no systematic effects on the net migration rate.

Appendix Figure A14: Event-study estimates for additional outcomes

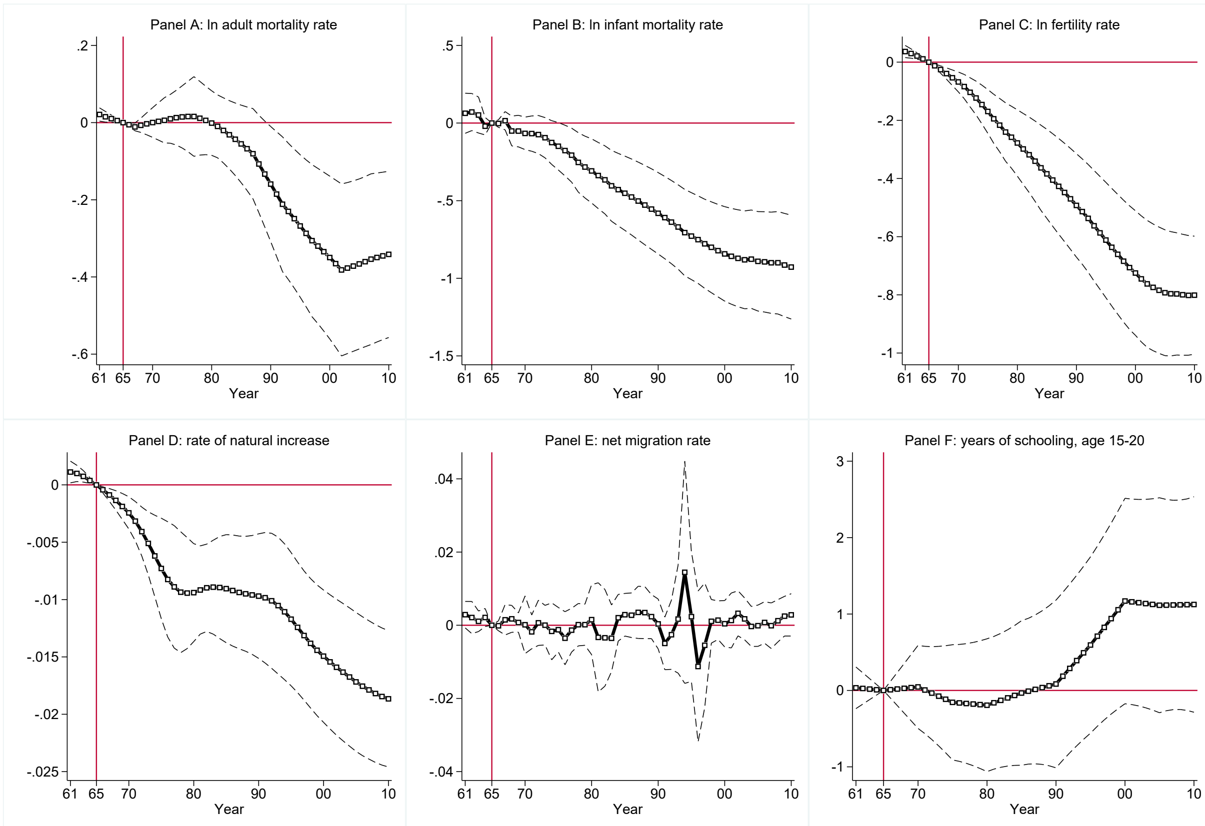

Notes: *This figure shows event-study estimates for the country-level outcomes: ln adult mortality rate (Panel A), ln infant mortality rate (Panel B), ln total fertility rate (Panel C), rate of natural population increase (Panel D), net migration rate (Panel E), years of schooling, age 15-20 (Panel F). These outcomes are also used in Table 6 of the main paper. The sample period is 1961-2010. All regressions include country and year fixed effects. The vertical red lines mark 1964, the last pre-Green Revolution period, which is also the omitted comparison year. The dashed lines indicate the 95-percent confidence bands. Standard errors are clustered at the country level.*

One might worry about possible negative pre-trends in fertility and the RNI, but they are hard to assess directly because of the limited length of the pre-Green Revolution period. But because our results for population size, which in the absence of any migration effects entirely is driven by RNI, becomes stronger when controlling for pre-trends, the same must be true for RNI. Moreover, in Appendix Figure A15, we demonstrate that our findings for fertility and RNI are robust to controlling for pre-Green Revolution linear trends estimated based on the short pre-period 1961-1964.

Appendix Figure A15: Event-study estimates for RNI and fertility w. pre-GR trends

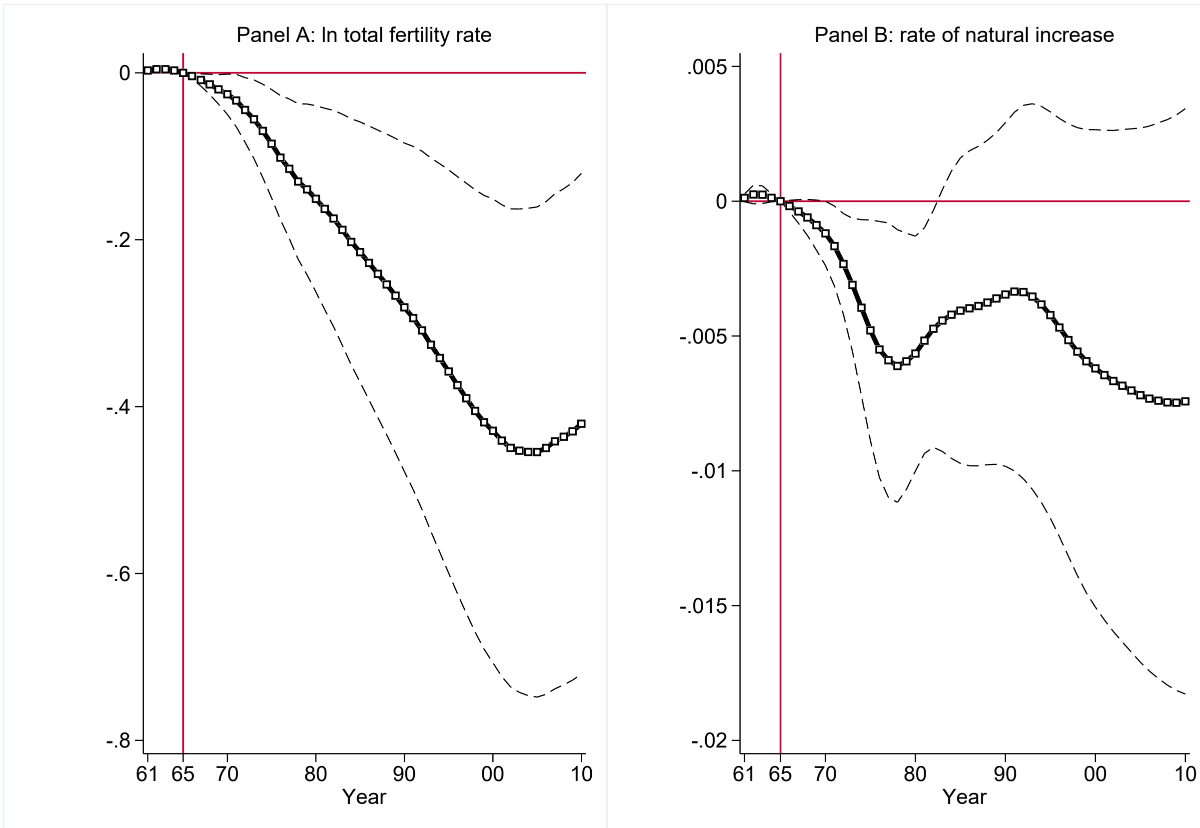

Notes: *This figure shows event-study estimates for the country-level outcomes: ln total fertility rate (Panel A) rate of natural population increase (Panel B). The sample period is 1961-2010 (due to data availability). All regressions include country and year fixed effects as well as pre-GR linear trends. The vertical line indicates the onset of the GR in 1965. The omitted comparison year is 1965. The dashed lines indicate the 95-percent confidence bands.*

Appendix Table A13 shows how the Green Revolution impacted other measures of economic development. We are here restricted in our choice of outcome variables by the fact that there only exists comprehensive pre-Green Revolution data for the developing world for few relevant outcomes. We find evidence suggesting that the supply of health care improved in terms of physicians per capita and hospital beds per capita (columns 1 and 2), although the results are not statistically significant in the case of the former. We also see evidence of increased up-take of non-agricultural technologies such as landline telephone subscriptions per capita (Column 3) and electricity consumption per capita (Column 4). The latter is a logged expenditure measure, and so electric power consumption has been increasing more than overall income (3.4 vs. 2.1). The point estimate suggests that development aid has decreased in countries more affected by the Green Revolution, consistent with these countries growing richer, but the estimate is statistically insignificant. We do not find any strong evidence that the Green Revolution led to increasing government consumption relative to GDP, but consistent with structural transformation, the Green Revolution seems to have increased the share of manufacturing in GDP (Column 7).

Appendix Table A13: Effects on measures of development

	Dependent variable:						
	ln physicians per capia	*ln hospital beds per capita*	*ln telephone subscripti*	*ln electricity consumption per*	*ln development aid per capita*	*government spending / GDP*	*Manufacturing GDP*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\ln \widehat{GR} \times initial\ AES$	1.103	2.162**	3.786**	5.833***	-2.461	0.666	0.295***
	(1.017)	(1.071)	(1.654)	(2.029)	(2.080)	(13.34)	(0.104)
Controls (x year FE)							
Initial AES	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,081	3,920	3,961	2,564	4,187	3,455	3,124
Countries	86	84	86	67	86	80	84

Notes: *This table reports effects on additional measures of economic development. The specific outcome is indicated in the top row. All specifications include country, year fixed effects, and the initial agricultural employment share interacted with year fixed effects. The sample period is 1961-2010. Standard errors are based on a two-step wild cluster restricted bootstrap procedure, which takes into account that* $\ln \widehat{GR}$ *is generated.*
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

I Effect sizes

This section outlines how we calculate the counterfactual scenarios discussed in Section 7 (and reported in Table 6) of the main paper. The counterfactual scenario with no Green Revolution is constructed as follows. First, we use the estimates reported in Column 2 of Table 3 (of the main paper) to calculate GDP per capita and population size in every sample country under the assumption that the predicted GR yields are zero in 2010. Second, we use the predicted population size without the Green Revolution to calculate counterfactual GDP levels. Third, we sum GDP levels and populations across the sample countries. Fourth, we divide the resulting aggregate GDP level by the aggregate population size to obtain a counterfactual aggregate GDP per capita level. We follow the same four steps when calculating counterfactual yields, except that we rely on the estimates reported in Column 1 of Table 3 in the main paper.

We construct the counterfactual scenarios with delays in the Green Revolution in the same way as we construct the counterfactual scenario without the Green Revolution. The only difference is that we lag the predicted GR yields by ten and 25 years, respectively. This means that in the counterfactual scenario with a ten year delay of the Green Revolution, for example, we assume that the contribution to yields from HYVs in year 2010 is equal to the actual contribution in year 2000.

The three counterfactual scenarios are calculated for the full sample, and, ind the case of GDP per capita, for various subsamples. The countries contained in the subsamples are listed in Appendix Table A4.

References

Barro, Robert J and Jong Wha Lee. 2013. "A new data set of educational attainment in the world, 1950–2010." *Journal of Development Economics* 104:184–198.

Bhuller, Manudeep, Tarjei Havnes, Edwin Leuven, and Magne Mogstad. 2013. "Broadband internet: An information superhighway to sex crime?" *Review of Economic Studies* 80 (4):1237–1266.

Buera, Francisco J, Alexander Monge-Naranjo, and Giorgio E Primiceri. 2011. "Learning the wealth of nations." *Econometrica* 79 (1):1–45.

Evenson, Robert E and Douglas Gollin. 2003a. "Assessing the impact of the Green Revolution, 1960 to 2000." *Science* 300 (5620):758–762.

———. 2003b. *Crop variety improvement and its effect on productivity: The impact of international agricultural research*. Cabi.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review* 105 (10):3150–3182. URL `http://www.ggdc.net/pwt/`.

Feyrer, James. 2009. "Trade and Income–Exploiting Time Series in Geography." *NBER Working Paper 14910* .

Gallup, John Luke and Jeffrey D Sachs. 2001. "The economic burden of malaria." *The American Journal of Tropical Nedicine and Hygiene* 64 (1_suppl):85–96.

Galor, Oded and Ömer Özak. 2016. "The agricultural origins of time preference." *American Economic Review* 106 (10):3064–3103.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2018. "Bartik instruments: What, when, why, and how." Tech. rep., National Bureau of Economic Research.

Gollin, Douglas, David Lagakos, and Michael E Waugh. 2014. "The Agricultural Productivity Gap." *The Quarterly Journal of Economics* 129 (2):939–993.

Gollin, Douglas, Stephen L Parente, and Richard Rogerson. 2007. "The food problem and the evolution of international income levels." *Journal of Monetary Economics* 54 (4):1230–1255.

Goodman-Bacon, Andrew. 2018. "Difference-in-differences with variation in treatment timing." Tech. rep., National Bureau of Economic Research.

Hicks, Joan Hamory, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2017. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." Tech. rep., National Bureau of Economic Research.

Inklaar, Robert, Harmen de Jong, Jutta Bolt, and Jan van Zanden. 2018. "Rebasing'Maddison': new income comparisons and the shape of long-run economic development." Tech. rep., Groningen Growth and Development Centre, University of Groningen.

Mankiw, N Gregory, David Romer, and David N Weil. 1992. "A contribution to the empirics of economic growth." *The quarterly journal of economics* 107 (2):407–437.

Mitchell, Brian R. 1982. "International historical statistics: Africa and Asia." .

Nunn, Nathan and Diego Puga. 2012. "Ruggedness: The blessing of bad geography in Africa." *Review of Economics and Statistics* 94 (1):20–36.

Rose, Beth. 1985. *Appendix to: The rice economy of Asia, Rice Statistics by Country*. The Johns Hopkins University Press.

Sachs, Jeffrey D, Andrew Warner, Anders Åslund, and Stanley Fischer. 1995. "Economic reform and the process of global integration." *Brookings papers on economic activity* 1995 (1):1–118.

Wingender, Asger Moll. 2014. "Structural transformation in the 20th Century: A new database on agricultural employment around the world." *Available at SSRN 2533360* .